



User Guide

AWS Certification



AWS Certification: User Guide

Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

What Is AWS Certification?	1
Overview of AWS Certification	1
AWS Certification Pathways	2
Overview of AWS Certification Pathways	2
Architecture Roles	2
Solutions Architect	2
Application Architect	3
Data Analytics Roles	3
Cloud Data Engineer	3
Development Roles	4
Software Development Engineer	4
Operations Roles	4
Systems Administrator	4
Cloud Engineer	4
DevOps Roles	5
Test Engineer	5
Cloud DevOps Engineer	5
DevSecOps Engineer	5
Security Roles	6
Cloud Security Engineer	6
Cloud Security Architect	6
Networking Roles	6
Network Engineer	6
AI/ML Roles	7
Prompt Engineer	7
Machine Learning Engineer	7
Machine Learning Ops Engineer	7
Data Scientist	8
Getting Started with Your Certification Path	8
Examens de niveau fondamental	9
AWS Certified Cloud Practitioner (CLF-C02)	9
Introduction	9
Description du candidat cible	10
Contenu de l'examen	11

Domaines de contenu	11
Services AWS pour l'examen	12
Domaine 1 : Concepts du cloud (24 % du contenu évalué)	12
Domaine 2 : Sécurité et conformité (30 % du contenu évalué)	14
Domaine 3 : Technologie et services cloud (34 % du contenu évalué)	17
Domaine 4 : Facturation, tarification et support (12 % du contenu évalué)	21
In-Scope AWS Services	23
Services AWS hors du champ d'application	29
AWS Certified AI Practitioner (AIF-C01)	34
Introduction	35
Description du candidat cible	35
Contenu de l'examen	36
Domaines de contenu	37
Services AWS pour l'examen	38
Domaine 1 : Principes fondamentaux de l'IA et du ML (20 % du contenu noté)	38
Domaine 2 : Principes fondamentaux de l'IA générative (24 % du contenu noté)	40
Domaine 3 : Applications des modèles de fondation (28 % du contenu noté)	41
Domaine 4 : Directives pour une IA responsable (14 % du contenu noté)	43
Domaine 5 : Sécurité, conformité et gouvernance pour les solutions d'IA (14 % du contenu noté)	45
Services AWS concernés	46
Services AWS non concernés	49
Associate Exams	56
AWS Certified Developer - Associate (DVA-C02)	56
Introduction	57
Target Candidate Description	57
Exam Content	58
Content Domains	58
AWS Services for the Exam	59
Domain 1: Development with Services (32% of the exam content)	59
Domain 2: Security (26% of the exam content)	61
Domain 3: Deployment (24% of the exam content)	63
Domain 4: Troubleshooting and Optimization (18% of the exam content)	66
In-Scope AWS Services	68
Out-of-Scope AWS Services	71
AWS Certified SysOps Administrator - Associate (SOA-C02)	74

Introduction	75
Target Candidate Description	75
Exam Content	76
Content Domains	76
AWS Services for the Exam	77
Domain 1: Monitoring, Logging, and Remediation (20% of the exam content)	77
Domain 2: Reliability and Business Continuity (16% of the exam content)	78
Domain 3: Deployment, Provisioning, and Automation (18% of the exam content)	79
Domain 4: Security and Compliance (16% of the exam content)	80
Domain 5: Networking and Content Delivery (18% of the exam content)	80
Domain 6: Cost and Performance Optimization (12% of the exam content)	81
In-Scope AWS Services	82
Out-of-Scope AWS Services	86
AWS Certified Solutions Architect - Associate (SAA-C03)	89
Introduction	89
Target Candidate Description	90
Exam Content	90
Content Domains	91
AWS Services for the Exam	91
Domain 1: Design Secure Architectures (30% of the exam content)	91
Domain 2: Design Resilient Architectures (26% of the exam content)	93
Domain 3: Design High-Performing Architectures (24% of the exam content)	96
Domain 4: Design Cost-Optimized Architectures (20% of the exam content)	99
In-Scope AWS Services	103
Out-of-Scope AWS Services	108
AWS Certified Machine Learning Engineer - Associate (MLA-C01)	111
Introduction	112
Target Candidate Description	112
Exam Content	113
Content Domains	114
AWS Services for the Exam	114
Domain 1: Data Preparation for Machine Learning (ML) (28% of the exam content)	115
Domain 2: ML Model Development (26% of the exam content)	117
Domain 3: Deployment and Orchestration of ML Workflows (22% of the exam content)	119
Domain 4: ML Solution Monitoring, Maintenance, and Security (24% of the exam content) ..	122
In-Scope AWS Services	124

AWS Certified Data Engineer - Associate (DEA-C01)	129
Introduction	130
Target Candidate Description	130
Exam Content	131
Content Domains	132
AWS Services for the Exam	132
Domain 1: Data Ingestion and Transformation (36% of the exam content)	132
Domain 2: Data Store Management (26% of the exam content)	135
Domain 3: Data Operations and Support (22% of the exam content)	138
Domain 4: Security and Governance (16% of the exam content)	140
In-Scope AWS Services	143
Professional Exams	148
AWS Certified Solutions Architect - Professional (SAP-C02)	148
Introduction	149
Target Candidate Description	149
Exam Content	149
Content Domains	150
AWS Services for the Exam	150
Domain 1: Design Solutions for Organizational Complexity (26% of the exam content)	151
Domain 2: Design for New Solutions (29% of the exam content)	153
Domain 3: Migration Planning (15% of the exam content)	156
Domain 4: Cost Control (10% of the exam content)	158
Domain 5: Continuous Improvement for Existing Solutions (20% of the exam content)	160
In-Scope AWS Services	162
Out-of-Scope AWS Services	168
AWS Certified DevOps Engineer - Professional (DOP-C02)	171
Introduction	172
Target Candidate Description	172
Exam Content	173
Content Domains	173
AWS Services for the Exam	173
Domain 1: SDLC Automation (22% of the exam content)	174
Domain 2: Configuration Management and IaC (17% of the exam content)	176
Domain 3: Resilient Cloud Solutions (15% of the exam content)	178
Domain 4: Monitoring and Logging (15% of the exam content)	180
Domain 5: Incident and Event Response (14% of the exam content)	182

Domain 6: Security and Compliance (17% of the exam content)	183
In-Scope AWS Services	185
Specialty Exams	189
AWS Certified Machine Learning - Specialty (MLS-C01)	189
Introduction	190
Target Candidate Description	190
Exam Content	191
Content Domains	191
AWS Services for the Exam	192
Domain 1: Data Engineering (20% of the exam content)	192
Domain 2: Exploratory Data Analysis (24% of the exam content)	193
Domain 3: Modeling (36% of the exam content)	194
Domain 4: Machine Learning Implementation and Operations (20% of the exam content)	196
In-Scope AWS Services	197
Out-of-Scope AWS Services	200
AWS Certified Security - Specialty (SCS-C02)	201
Introduction	201
Target Candidate Description	202
Exam Content	203
Content Domains	203
AWS Services for the Exam	204
Domain 1: Threat Detection and Incident Response (20% of the exam content)	204
Domain 2: Security Logging and Monitoring (18% of the exam content)	206
Domain 3: Infrastructure Security (16% of the exam content)	209
Domain 4: Identity and Access Management (20% of the exam content)	212
Domain 5: Data Protection (18% of the exam content)	214
Domain 6: Management and Security Governance (8% of the exam content)	216
In-Scope AWS Services	219
AWS Certified Advanced Networking - Specialty (ANS-C01)	222
Introduction	223
Target Candidate Description	223
Exam Content	223
Content Domains	224
AWS Services for the Exam	224
Domain 1: Network Design (30% of the exam content)	224
Domain 2: Network Implementation (26% of the exam content)	228

Domain 3: Network Management and Operation (20% of the exam content)

232

Domain 4: Network Security, Compliance, and Governance (24% of the exam content)

234

In-Scope AWS Services

237

Out-of-Scope AWS Services

240

Document history

243

What Is AWS Certification?

Overview of AWS Certification

AWS Certification(s) help individuals validate their knowledge, skills, and expertise, and help organizations identify and hire the right talent. Today, we offer certifications across three levels of proficiency: Foundational (validating fundamental knowledge of AWS Cloud), Associate and Professional (validating technical skills and expertise), as well as Specialty (showcasing expertise in specific areas of focus) certifications. [Learn more](#).

AWS Certification Pathways

Rubriques

- [Overview of AWS Certification Pathways](#)
- [Architecture Roles](#)
- [Data Analytics Roles](#)
- [Development Roles](#)
- [Operations Roles](#)
- [DevOps Roles](#)
- [Security Roles](#)
- [Networking Roles](#)
- [AI/ML Roles](#)
- [Getting Started with Your Certification Path](#)

Overview of AWS Certification Pathways

AWS offers recommended certification paths aligned with top cloud job roles. These pathways are designed to help you achieve your career goals by earning the most relevant certifications for your role. Note that these are recommended pathways, and you are not required to follow them exactly.

For IT professionals with 1-3 years of IT or STEM background, it's recommended to start with an Associate-level AWS Certification that aligns with your role. AWS Certified AI Practitioner is also recommended to validate conceptual AI knowledge.

For those new to IT and cloud or from non-IT backgrounds, it's recommended to start with AWS Certified Cloud Practitioner to validate foundational AWS Cloud knowledge, then earn AWS Certified AI Practitioner to showcase AI knowledge.

Architecture Roles

Solutions Architect

Solutions Architects design, develop, and manage cloud infrastructure and assets, and work with DevOps to migrate applications to the cloud.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified Solutions Architect - Associate
3. AWS Certified AI Practitioner (Recommended for IT/cloud professionals to leverage AI)
4. AWS Certified Solutions Architect - Professional

Application Architect

Application Architects design significant aspects of application architecture including user interface, middleware, and infrastructure, and ensure enterprise-wide scalable, reliable, and manageable systems.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified Solutions Architect - Associate
3. AWS Certified AI Practitioner (Recommended for IT/cloud professionals to leverage AI)
4. AWS Certified Solutions Architect - Professional

Data Analytics Roles

Cloud Data Engineer

Cloud Data Engineers automate collection and processing of structured/semi-structured data and monitor data pipeline performance.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified Data Engineer - Associate
3. AWS Certified Machine Learning - Specialty (Recommended for IT/Cloud professionals working on AI/ML projects)

Development Roles

Software Development Engineer

Software Development Engineers develop, construct, and maintain software across platforms and devices.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified Developer - Associate
3. AWS Certified AI Practitioner (Recommended for IT/cloud professionals to leverage AI)

Operations Roles

Systems Administrator

Systems Administrators install, upgrade, and maintain computer components and software, and integrate automation processes.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified SysOps Administrator - Associate

Cloud Engineer

Cloud Engineers implement and operate an organization's networked computing infrastructure and implement security systems to maintain data safety.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified SysOps Administrator - Associate

DevOps Roles

Test Engineer

Test Engineers embed testing and quality best practices for software development from design to release, throughout the product life cycle.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified Developer - Associate

Cloud DevOps Engineer

Cloud DevOps Engineers design, deploy, and operate large-scale global hybrid cloud computing environments, advocating for end-to-end automated CI/CD DevOps pipelines.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified Developer - Associate
3. AWS Certified DevOps Engineer - Professional
4. AWS Certified Machine Learning - Specialty (Optional)

DevSecOps Engineer

DevSecOps Engineers accelerate enterprise cloud adoption while enabling rapid and stable delivery of capabilities using CI/CD principles, methodologies, and technologies.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified Developer - Associate
3. AWS Certified Security - Specialty
4. AWS Certified DevOps Engineer - Professional
5. AWS Certified Machine Learning - Specialty (Recommended for IT/Cloud professionals working on AI/ML projects)

Security Roles

Cloud Security Engineer

Cloud Security Engineers design computer security architecture and develop detailed cyber security designs. They develop, execute, and track performance of security measures to protect information.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified Security - Specialty
3. AWS Certified AI Practitioner (Recommended for IT/cloud professionals to secure AI/ML systems)

Cloud Security Architect

Cloud Security Architects design and implement enterprise cloud solutions applying governance to identify, communicate, and minimize business and technical risks.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified Solutions Architect - Associate
3. AWS Certified Security - Specialty
4. AWS Certified AI Practitioner (Recommended for IT/cloud professionals to secure AI/ML systems)

Networking Roles

Network Engineer

Network Engineers design and implement computer and information networks, such as local area networks (LAN), wide area networks (WAN), intranets, extranets, etc.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified Solutions Architect - Associate
3. AWS Certified Advanced Networking - Specialty

AI/ML Roles

Prompt Engineer

Prompt Engineers design, test, and refine text prompts to optimize the performance of AI language models.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified AI Practitioner
3. AWS Certified Machine Learning - Specialty

Machine Learning Engineer

Machine Learning Engineers research, build, and design artificial intelligence (AI) systems to automate predictive models, and design machine learning systems, models, and schemes.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified Machine Learning Engineer - Associate
3. AWS Certified AI Practitioner (Optional for AI/ML professionals)
4. AWS Certified Machine Learning - Specialty

Machine Learning Ops Engineer

Machine Learning Ops Engineers build and maintain AI and ML platforms and infrastructure. They design, implement, and operationally support AI/ML model activity and deployment infrastructure.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified Machine Learning Engineer - Associate
3. AWS Certified DevOps Engineer - Professional
4. AWS Certified AI Practitioner (Optional for AI/ML professionals)

5. AWS Certified Machine Learning - Specialty

Data Scientist

Data Scientists develop and implement AI/ML models to solve business problems. They train and fine tune models and evaluate their performance.

Recommended Certification Path:

1. AWS Certified Cloud Practitioner (Optional for IT/cloud professionals)
2. AWS Certified Data Engineer - Associate
3. AWS Certified AI Practitioner (Optional for AI/ML professionals)
4. AWS Certified Machine Learning - Specialty

Getting Started with Your Certification Path

Ready to get started? Select the exam(s) aligned with your career goals and follow these steps to approach exam day with confidence:

1. Get to know the exam with exam-style questions
2. Refresh your AWS knowledge and skills with AWS Skill Builder
3. Review and practice for your exam in AWS Skill Builder
4. Assess your exam readiness and approach exam day with confidence

Examens de niveau fondamental

Les examens de certification AWS de niveau fondamental sont conçus pour les personnes qui découvrent AWS et le cloud computing.

Rubriques

- [AWS Certified Cloud Practitioner \(CLF-C02\)](#)
- [AWS Certified AI Practitioner \(AIF-C01\)](#)

AWS Certified Cloud Practitioner (CLF-C02)

L'examen AWS Certified Cloud Practitioner (CLF-C02) est destiné aux personnes qui peuvent démontrer efficacement une connaissance globale du Cloud AWS, indépendamment d'un rôle professionnel spécifique.

Rubriques

- [Introduction](#)
- [Description du candidat cible](#)
- [Contenu de l'examen](#)
- [Domaines de contenu](#)
- [Services AWS pour l'examen](#)
- [Domaine 1 : Concepts du cloud \(24 % du contenu évalué\)](#)
- [Domaine 2 : Sécurité et conformité \(30 % du contenu évalué\)](#)
- [Domaine 3 : Technologie et services cloud \(34 % du contenu évalué\)](#)
- [Domaine 4 : Facturation, tarification et support \(12 % du contenu évalué\)](#)
- [In-Scope AWS Services](#)
- [Services AWS hors du champ d'application](#)

Introduction

L'examen AWS Certified Cloud Practitioner (CLF-C02) est destiné aux personnes qui peuvent démontrer efficacement une connaissance globale du Cloud AWS, indépendamment d'un rôle professionnel spécifique.

L'examen valide également la capacité d'un candidat à effectuer les tâches suivantes :

- Expliquer la valeur du Cloud AWS
- Comprendre et expliquer le modèle de responsabilité partagée AWS
- Comprendre le cadre AWS Well-Architected
- Comprendre les bonnes pratiques en matière de sécurité
- Comprendre les coûts, les aspects économiques et les pratiques de facturation du Cloud AWS
- Décrire et positionner les principaux services AWS, y compris les services de calcul, de réseau, de base de données et de stockage
- Identifier les services AWS pour les cas d'utilisation courants

Description du candidat cible

Le candidat cible a jusqu'à six mois d'expérience en matière de conception, de mise en œuvre et/ou d'opérations du Cloud AWS. Ce candidat n'en est peut-être qu'aux premiers stades de sa carrière dans le Cloud AWS ou peut travailler avec des personnes occupant des rôles dans le Cloud AWS.

Le candidat cible doit disposer de connaissances AWS dans les domaines suivants :

- Concepts relatifs au Cloud AWS
- Sécurité et conformité dans le Cloud AWS
- Principaux services AWS
- Économie du Cloud AWS

La liste suivante contient les tâches que le candidat cible n'est pas censé être en mesure d'effectuer. Cette liste n'est pas exhaustive. Ces tâches sortent du champ d'application de l'examen :

- Codage
- Conception d'architectures cloud
- Résolution des problèmes
- Implémentation
- Tests de charge et de performance

Contenu de l'examen

L'examen comporte deux types de questions :

- Choix multiple : une réponse correcte et trois réponses incorrectes (réponses piège)
- Réponse multiple : plusieurs réponses correctes sur cinq options de réponse ou plus

Sélectionnez une ou plusieurs réponses qui complètent correctement la déclaration ou répondent à la question. Les réponses piège, ou réponses incorrectes, sont des options qu'un candidat ayant des connaissances ou compétences incomplètes est susceptible de choisir. Les réponses piège sont généralement des réponses plausibles qui correspondent au contenu.

Les questions sans réponse sont notées comme incorrectes ; aucune pénalité n'est appliquée lorsque le candidat devine une réponse. L'examen comporte 50 questions qui ont une incidence sur votre score.

L'examen comporte 15 questions non notées qui n'ont pas d'incidence sur votre score. AWS recueille des informations sur les performances dans le cadre de ces questions non notées. L'objectif est d'évaluer ces questions en vue d'une utilisation ultérieure en tant que questions notées. Ces questions non notées ne sont pas identifiées comme telles dans l'examen.

L'examen AWS Certified Cloud Practitioner (CLF-C02) est un examen de type « réussite ou échec ». L'examen est évalué selon une norme minimale établie par les professionnels d'AWS. Ces derniers observent les bonnes pratiques et directives en matière de certification.

Vos résultats à l'examen sont présentés sous la forme d'un score gradué de 100 à 1 000. Le score minimal pour réussir est de 700. Votre score indique vos performances lors de l'examen dans son ensemble et si vous avez réussi ce dernier. Les modèles de notation graduée permettent de mettre en correspondance des scores de différents formulaires d'examen qui peuvent présenter des niveaux de difficulté légèrement différents.

Domaines de contenu

L'examen comporte les domaines de contenu et les pondérations suivants :

- Domaine 1 : Concepts du cloud (24 % du contenu évalué)
- Domaine 2 : Sécurité et conformité (30 % du contenu évalué)
- Domaine 3 : Technologie et services cloud (34 % du contenu évalué)
- Domaine 4 : Facturation, tarification et support (12 % du contenu évalué)

Pour des informations détaillées sur chaque domaine, consultez les sections suivantes :

- [the section called “Domaine 1 : Concepts du cloud \(24 % du contenu évalué\)”](#)
- [the section called “Domaine 2 : Sécurité et conformité \(30 % du contenu évalué\)”](#)
- [the section called “Domaine 3 : Technologie et services cloud \(34 % du contenu évalué\)”](#)
- [the section called “Domaine 4 : Facturation, tarification et support \(12 % du contenu évalué\)”](#)

Services AWS pour l'examen

L'examen AWS Certified Cloud Practitioner couvre des services AWS spécifiques qui sont pertinents pour les fondamentaux du cloud. Comprendre quels services sont dans le champ d'application et lesquels sont hors du champ d'application peut vous aider à concentrer vos efforts de préparation.

Pour des informations détaillées sur les services AWS couverts dans l'examen, consultez les sections suivantes :

- [the section called “In-Scope AWS Services”](#)
- [the section called “Services AWS hors du champ d'application”](#)

Domaine 1 : Concepts du cloud (24 % du contenu évalué)

Le Domaine 1 couvre les Concepts du cloud et représente 24 % du contenu évalué de l'examen.

Rubriques

- [Énoncé de tâche 1.1 : définir les avantages du Cloud AWS](#)
- [Énoncé de tâche 1.2 : identifier les principes de conception du Cloud AWS](#)
- [Énoncé de tâche 1.3 : comprendre les avantages et les stratégies de la migration vers le Cloud AWS](#)
- [Énoncé de tâche 1.4 : comprendre les concepts de l'économie du cloud](#)

Énoncé de tâche 1.1 : définir les avantages du Cloud AWS

Connaissance des éléments suivants :

- Proposition de valeur du Cloud AWS

Compétences dans les domaines suivants :

- Comprendre les avantages d'une infrastructure globale (par exemple, la rapidité du déploiement, la portée mondiale)
- Comprendre les avantages de la haute disponibilité, de l'élasticité et de l'agilité

Énoncé de tâche 1.2 : identifier les principes de conception du Cloud AWS

Connaissance des éléments suivants :

- Cadre AWS Well-Architected

Compétences dans les domaines suivants :

- Comprendre les piliers du cadre Well-Architected (par exemple, l'excellence opérationnelle, la sécurité, la fiabilité, l'efficacité des performances, l'optimisation des coûts, la durabilité)
- Identifier les différences entre les piliers du cadre Well-Architected

Énoncé de tâche 1.3 : comprendre les avantages et les stratégies de la migration vers le Cloud AWS

Connaissance des éléments suivants :

- Stratégies d'adoption du cloud
- Ressources pour soutenir le parcours de migration vers le cloud

Compétences dans les domaines suivants :

- Comprendre les composants du Framework d'adoption du Cloud AWS (AWS CAF). (Par exemple, réduction des risques métier ; amélioration des performances environnementales, sociales et de gouvernance [ESG] ; augmentation des revenus ; augmentation de l'efficacité opérationnelle)
- Identifier les stratégies de migration appropriées (par exemple, réplication des bases de données, utilisation d'AWS Snowball)

Énoncé de tâche 1.4 : comprendre les concepts de l'économie du cloud

Connaissance des éléments suivants :

- Aspects de l'économie du cloud
- Économies de coûts liées au passage au cloud

Compétences dans les domaines suivants :

- Comprendre le rôle des coûts fixes par rapport à celui des coûts variables
- Comprendre les coûts associés aux environnements sur site
- Comprendre les différences entre les stratégies d'octroi de licences (par exemple, le modèle BYOL [Bring Your Own License] comparé aux licences incluses)
- Comprendre le concept de dimensionnement
- Identifier les avantages de l'automatisation
- Comprendre les économies d'échelle (par exemple, les réductions de coûts)

Domaine 2 : Sécurité et conformité (30 % du contenu évalué)

Le Domaine 2 couvre la Sécurité et la conformité et représente 30 % du contenu évalué de l'examen.

Rubriques

- [Énoncé de tâche 2.1 : comprendre le modèle de responsabilité partagée AWS](#)
- [Énoncé de tâche 2.2 : comprendre les concepts de sécurité, de gouvernance et de conformité du Cloud AWS](#)
- [Énoncé de tâche 2.3 : identifier les capacités de gestion des accès dans AWS](#)
- [Énoncé de tâche 2.4 : identifier les composants et les ressources pour la sécurité](#)

Énoncé de tâche 2.1 : comprendre le modèle de responsabilité partagée AWS

Connaissance des éléments suivants :

- Modèle de responsabilité partagée AWS

Compétences dans les domaines suivants :

- Reconnaître les composantes du modèle de responsabilité partagée AWS
- Décrire les responsabilités du client sur AWS
- Décrire les responsabilités d'AWS
- Décrire les responsabilités partagées du client et d'AWS
- Décrire comment les responsabilités d'AWS et celles des clients peuvent évoluer, en fonction du service utilisé (par exemple, Amazon RDS, AWS Lambda, Amazon EC2)

Énoncé de tâche 2.2 : comprendre les concepts de sécurité, de gouvernance et de conformité du Cloud AWS

Connaissance des éléments suivants :

- Concepts de conformité et de gouvernance AWS
- Avantages de la sécurité du cloud (par exemple, le chiffrement)
- Où capturer et localiser les journaux qui sont associés à la sécurité du cloud

Compétences dans les domaines suivants :

- Identifier où trouver les informations relatives à la conformité AWS (par exemple, AWS Artifact)
- Comprendre les besoins en matière de conformité selon les emplacements géographiques ou les secteurs d'activité (par exemple, la conformité AWS)
- Décrire comment les clients protègent les ressources sur AWS (par exemple, Amazon Inspector, AWS Security Hub, Amazon GuardDuty, AWS Shield)
- Identifier les options de chiffrement (par exemple, chiffrement en transit, chiffrement au repos)
- Reconnaître les services qui contribuent à la gouvernance et à la conformité (par exemple, la surveillance avec Amazon CloudWatch, l'audit avec AWS CloudTrail, AWS Audit Manager et AWS Config, la création de rapports avec les rapports d'accès)
- Reconnaître les exigences de conformité qui varient selon les services AWS

Énoncé de tâche 2.3 : identifier les capacités de gestion des accès dans AWS

Connaissance des éléments suivants :

- Gérer les identités et les accès (par exemple, Gestion des identités et des accès AWS [AWS IAM])

- Importance de la protection du compte utilisateur racine AWS
- Principe du moindre privilège
- AWS IAM Identity Center (AWS Single Sign-On)

Compétences dans les domaines suivants :

- Comprendre les clés d'accès, les politiques de mots de passe et le stockage des informations d'identification (par exemple, AWS Secrets Manager, AWS Systems Manager)
- Identifier les méthodes d'authentification dans AWS (par exemple, l'authentification multifacteur [MFA], IAM Identity Center, les rôles IAM intercomptes)
- Définir les groupes, les utilisateurs, les politiques personnalisées et les politiques gérées en respectant le principe du moindre privilège
- Identifier les tâches que seul l'utilisateur racine du compte peut effectuer
- Comprendre les méthodes permettant de protéger l'utilisateur racine
- Comprendre les types de gestion d'identité (par exemple, fédérée)

Énoncé de tâche 2.4 : identifier les composants et les ressources pour la sécurité

Connaissance des éléments suivants :

- Fonctions de sécurité proposées par AWS
- Documentation relative à la sécurité fournie par AWS

Compétences dans les domaines suivants :

- Décrire les fonctionnalités et services de sécurité AWS (par exemple, AWS WAF, AWS Firewall Manager, AWS Shield, Amazon GuardDuty)
- Comprendre que des produits de sécurité tiers sont disponibles sur AWS Marketplace
- Identifier les sources d'information sur la sécurité AWS (par exemple, AWS Knowledge Center, Centre de sécurité AWS, Blog AWS sur la sécurité)
- Comprendre l'utilisation des services AWS pour identifier les problèmes de sécurité (par exemple, AWS Trusted Advisor)

Domaine 3 : Technologie et services cloud (34 % du contenu évalué)

Le Domaine 3 couvre la Technologie et les services cloud et représente 34 % du contenu évalué de l'examen.

Rubriques

- [Énoncé de tâche 3.1 : définir les méthodes de déploiement et d'opération dans le Cloud AWS](#)
- [Énoncé de tâche 3.2 : décrire l'infrastructure globale AWS](#)
- [Énoncé de tâche 3.3 : identifier les services de calcul d'AWS](#)
- [Énoncé de tâche 3.4 : identifier les services de base de données AWS](#)
- [Énoncé de tâche 3.5 : identifier les services de réseau AWS](#)
- [Task Statement 3.6: Identify AWS storage services](#)
- [Task Statement 3.7: Identify AWS artificial intelligence and machine learning \(AI/ML\) services and analytics services](#)
- [Task Statement 3.8: Identify services from other in-scope AWS service categories](#)

Énoncé de tâche 3.1 : définir les méthodes de déploiement et d'opération dans le Cloud AWS

Connaissance des éléments suivants :

- Diverses manières de mettre en service et d'utiliser des ressources dans le Cloud AWS
- Diverses manières d'accéder aux services AWS
- Types de modèles de déploiement de cloud

Compétences dans les domaines suivants :

- Choisir entre des options telles que l'accès programmatique (par exemple, API, SDK, CLI), la Console de gestion AWS et l'Infrastructure as Code (IaC)
- Évaluer les besoins pour déterminer s'il est nécessaire d'utiliser des opérations ponctuelles ou des processus reproductibles
- Identifier les modèles de déploiement (par exemple, cloud, hybride, sur site)

Énoncé de tâche 3.2 : décrire l'infrastructure globale AWS

Connaissance des éléments suivants :

- Régions, zones de disponibilité (AZ) et emplacements périphériques AWS
- Haute disponibilité
- Utiliser plusieurs régions
- Avantages des emplacements périphériques

Compétences dans les domaines suivants :

- Décrire les relations entre les régions, les zones de disponibilité (AZ) et les emplacements périphériques
- Décrire comment obtenir une haute disponibilité en utilisant plusieurs zones de disponibilité (AZ)
- Reconnaître que les zones de disponibilité (AZ) ne partagent pas de points de défaillance uniques
- Décrire quand utiliser plusieurs régions (par exemple, reprise après sinistre, continuité d'activités, faible latence pour les utilisateurs finaux, souveraineté des données)

Énoncé de tâche 3.3 : identifier les services de calcul d'AWS

Connaissance des éléments suivants :

- Services de calcul AWS

Compétences dans les domaines suivants :

- Reconnaître l'utilisation appropriée des différents types d'instance EC2 (par exemple, optimisé pour le calcul, optimisé pour le stockage)
- Reconnaître l'utilisation appropriée de diverses options de conteneurs (par exemple, Amazon ECS, Amazon EKS)
- Reconnaître l'utilisation appropriée des diverses options de calcul sans serveur (par exemple, AWS Fargate, Lambda)
- Reconnaître que le processus de mise à l'échelle automatique apporte de l'élasticité
- Identifier les objectifs des équilibres de charge

Énoncé de tâche 3.4 : identifier les services de base de données AWS

Connaissance des éléments suivants :

- Services de base de données AWS
- Migration de base de données

Compétences dans les domaines suivants :

- Décider quand utiliser des bases de données hébergées par EC2 ou des bases de données gérées par AWS
- Identifier les bases de données relationnelles (par exemple, Amazon RDS, Amazon Aurora)
- Identifier les bases de données NoSQL (par exemple, DynamoDB)
- Identifier les bases de données basées sur la mémoire (par exemple, Amazon ElastiCache)
- Identifier les outils de migration de bases de données (par exemple, AWS Database Migration Service [AWS DMS], AWS Schema Conversion Tool [AWS SCT])

Énoncé de tâche 3.5 : identifier les services de réseau AWS

Connaissance des éléments suivants :

- Services de réseau AWS

Compétences dans les domaines suivants :

- Identifier les composants d'un VPC (par exemple, sous-réseaux, passerelles)
- Comprendre la sécurité dans un VPC (par exemple, ACL réseau, groupes de sécurité, Amazon Inspector)
- Comprendre l'utilité d'Amazon Route 53
- Identifier les options de connectivité réseau vers AWS (par exemple AWS VPN, AWS Direct Connect)

Task Statement 3.6: Identify AWS storage services

Knowledge of:

- AWS storage services

Skills in:

- Identifying the uses for object storage
- Recognizing the differences in Amazon S3 storage classes
- Identifying block storage solutions
- Identifying file services
- Identifying cached file systems
- Understanding use cases for lifecycle policies
- Understanding use cases for AWS Backup

Task Statement 3.7: Identify AWS artificial intelligence and machine learning (AI/ML) services and analytics services

Knowledge of:

- AWS AI/ML services
- AWS analytics services

Skills in:

- Understanding AI/ML services and the tasks that they accomplish
- Identifying the services for data analytics

Task Statement 3.8: Identify services from other in-scope AWS service categories

Knowledge of:

- Application integration services
- Business application services
- Customer enablement services
- Developer tool services and capabilities
- End-user computing services

- Frontend web and mobile services
- IoT services

Skills in:

- Choosing the appropriate service to deliver messages and to send alerts and notifications
- Choosing the appropriate service to meet business application needs
- Choosing the appropriate option for business support assistance
- Identifying the tools to develop, deploy, and troubleshoot applications
- Identifying the services that can present the output of virtual machines (VMs) on end-user machines
- Identifying the services that can create and deploy frontend and mobile services
- Identifying the services that manage IoT devices

Domaine 4 : Facturation, tarification et support (12 % du contenu évalué)

Le Domaine 4 couvre la Facturation, la tarification et le support et représente 12 % du contenu évalué de l'examen.

Rubriques

- [Énoncé de tâche 4.1 : comparer les modèles de tarification AWS](#)
- [Énoncé de tâche 4.2 : comprendre les ressources pour la gestion de la facturation, du budget et des coûts](#)
- [Énoncé de tâche 4.3 : identifier les ressources techniques d'AWS et les options d'AWS Support](#)

Énoncé de tâche 4.1 : comparer les modèles de tarification AWS

Connaissance des éléments suivants :

- Options d'achat de capacités de calcul (par exemple, instances à la demande, instances réservées, instances Spot, Savings Plans, hôtes dédiés, instances dédiées, Réservations de capacité)
- Options et niveaux de stockage

Compétences dans les domaines suivants :

- Identifier quand utiliser les différentes options d'achat de capacités de calcul
- Décrire la flexibilité de l'instance réservée
- Décrire le comportement des instances réservées dans AWS Organizations
- Comprendre les coûts de transfert des données entrantes et ceux des données sortantes (par exemple, d'une région à une autre, au sein d'une même région)
- Comprendre les options de tarification pour les diverses options et différents niveaux de stockage

Énoncé de tâche 4.2 : comprendre les ressources pour la gestion de la facturation, du budget et des coûts

Connaissance des éléments suivants :

- Assistance et informations sur la facturation
- Informations sur la tarification des services AWS
- AWS Organizations
- Balises d'allocation des coûts AWS

Compétences dans les domaines suivants :

- Comprendre les utilisations appropriées et les fonctions d'AWS Budgets et d'AWS Cost Explorer
- Comprendre les utilisations appropriées et les fonctions du Calculateur de prix AWS
- Comprendre la facturation consolidée et l'allocation des coûts d'AWS Organizations
- Comprendre les différents types de balises de répartition des coûts et leur relation avec les rapports de facturation (par exemple, le Rapport d'utilisation et de coût AWS)

Énoncé de tâche 4.3 : identifier les ressources techniques d'AWS et les options d'AWS Support

Connaissance des éléments suivants :

- Ressources et documentation disponibles sur les sites web officiels d'AWS
- Programmes AWS Support
- Rôle du Réseau de partenaires AWS, y compris les fournisseurs indépendants de logiciels et les intégrateurs système

- AWS Support Center

Compétences dans les domaines suivants :

- Trouver des livres blancs, des blogs et de la documentation sur les sites web officiels d'AWS
- Identifier et localiser les ressources techniques AWS (par exemple, les Recommandations AWS, AWS Knowledge Center, AWS re:Post)
- Identifier les options d'AWS Support pour les clients AWS (par exemple, service clientèle et communautés, AWS Developer Support, AWS Business Support, AWS Enterprise On-Ramp Support, AWS Enterprise Support)
- Identifier le rôle de Trusted Advisor, du Tableau de bord AWS Health et de l'API AWS Health pour aider à gérer et surveiller les environnements en vue de l'optimisation des coûts
- Identifier le rôle de l'équipe AWS Trust and Safety pour signaler les abus des ressources AWS
- Comprendre le rôle des Partenaires AWS (par exemple AWS Marketplace, fournisseurs indépendants de logiciels, intégrateurs système)
- Identifier les avantages d'être un Partenaire AWS (par exemple, la formation et la certification des partenaires, les événements organisés par les partenaires, les remises sur volume des partenaires)
- Identifier les services clés offerts par AWS Marketplace (par exemple, la gestion des coûts, la gouvernance et les autorisations)
- Identifier les options d'assistance technique disponibles chez AWS (par exemple, AWS Professional Services, AWS Solutions Architects)

In-Scope AWS Services

The following list contains AWS services and features that are in scope for the AWS Certified Cloud Practitioner (CLF-C02) exam. This list is non-exhaustive and is subject to change. AWS offerings appear in categories that align with the offerings' primary functions.

Rubriques

- [Analytique](#)
- [Intégration d'applications](#)
- [Applications métier](#)
- [Gestion financière du cloud](#)
- [Calcul](#)

- [Conteneurs](#)
- [Émancipation du client](#)
- [Base de données](#)
- [Outils pour développeur](#)
- [Informatique pour l'utilisateur final](#)
- [Applications web et mobiles frontales](#)
- [Internet des objets \(IoT\)](#)
- [Machine learning](#)
- [Management et gouvernance](#)
- [Migration et transfert](#)
- [Mise en réseau et diffusion de contenu](#)
- [Sécurité, identité et conformité](#)
- [Sans serveur](#)
- [Stockage](#)

Analytique

- Amazon Athena
- Amazon EMR
- AWS Glue
- Amazon Kinesis
- Amazon OpenSearch Service
- Amazon QuickSight
- Amazon Redshift

Intégration d'applications

- Amazon EventBridge
- Amazon Simple Notification Service (Amazon SNS)
- Amazon Simple Queue Service (Amazon SQS)
- AWS Step Functions

Applications métier

- Amazon Connect
- Amazon Simple Email Service (Amazon SES)

Gestion financière du cloud

- AWS Budgets
- Rapports d'utilisation et de coût AWS
- AWS Cost Explorer
- AWS Marketplace

Calcul

- AWS Batch
- Amazon EC2
- AWS Elastic Beanstalk
- Amazon Lightsail
- AWS Outposts

Conteneurs

- Amazon Elastic Container Registry (Amazon ECR)
- Amazon Elastic Container Service (Amazon ECS)
- Amazon Elastic Kubernetes Service (Amazon EKS)

Émancipation du client

- AWS Support

Base de données

- Amazon Aurora
- Amazon DocumentDB

- Amazon DynamoDB
- Amazon ElastiCache
- Amazon Neptune
- Amazon RDS

Outils pour développeur

- AWS CLI
- AWS CodeBuild
- AWS CodePipeline
- AWS X-Ray

Informatique pour l'utilisateur final

- Amazon AppStream 2.0
- Amazon WorkSpaces
- Amazon WorkSpaces Secure Browser

Applications web et mobiles frontales

- AWS Amplify
- AWS AppSync

Internet des objets (IoT)

- AWS IoT Core

Machine learning

- Amazon Comprehend
- Amazon Kendra
- Amazon Lex
- Amazon Polly

- Amazon Q
- Amazon Rekognition
- IA Amazon SageMaker
- Amazon Textract
- Amazon Transcribe
- Amazon Translate

Management et gouvernance

- AWS Auto Scaling
- AWS CloudFormation
- AWS CloudTrail
- Amazon CloudWatch
- Optimiseur de calcul AWS
- AWS Config
- AWS Control Tower
- Tableau de bord AWS Health
- AWS License Manager
- Console de gestion AWS
- AWS Organizations
- AWS Service Catalog
- Service Quotas
- AWS Systems Manager
- AWS Trusted Advisor
- Outil AWS Well-Architected

Migration et transfert

- AWS Application Discovery Service
- AWS Application Migration Service

- AWS Database Migration Service (AWS DMS)
- Migration Evaluator
- AWS Migration Hub
- Outil de conversion des schémas AWS (AWS SCT)
- Gamme AWS Snow

Mise en réseau et diffusion de contenu

- Amazon API Gateway
- Amazon CloudFront
- AWS Direct Connect
- AWS Global Accelerator
- AWS PrivateLink
- Amazon Route 53
- AWS Transit Gateway
- Amazon VPC
- AWS VPN
- AWS Site-to-Site VPN
- AWS Client VPN

Sécurité, identité et conformité

- AWS Artifact
- AWS Audit Manager
- AWS Certificate Manager (ACM)
- AWS CloudHSM
- Amazon Cognito
- Amazon Detective
- AWS Directory Service
- AWS Firewall Manager
- Amazon GuardDuty

- Gestion des identités et des accès AWS (AWS IAM)
- AWS IAM Identity Center
- Amazon Inspector
- AWS Key Management Service (AWS KMS)
- Amazon Macie
- AWS Resource Access Manager (AWS RAM)
- AWS Secrets Manager
- AWS Security Hub
- AWS Shield
- AWS WAF

Sans serveur

- AWS Fargate
- AWS Lambda

Stockage

- AWS Backup
- Amazon Elastic Block Store (Amazon EBS)
- Amazon Elastic File System (Amazon EFS)
- AWS Elastic Disaster Recovery
- Amazon FSx
- Amazon S3
- Amazon S3 Glacier
- AWS Storage Gateway

Services AWS hors du champ d'application

La liste suivante contient les services et fonctionnalités AWS qui sont hors du champ d'application de l'examen AWS Certified Cloud Practitioner (CLF-C02). Cette liste n'est pas exhaustive et peut faire l'objet de modifications.

Rubriques

- [Analytique](#)
- [Intégration d'applications](#)
- [Applications métier](#)
- [Calcul](#)
- [Gestion des coûts](#)
- [Émancipation du client](#)
- [Gestion financière du cloud](#)
- [Base de données](#)
- [Outils pour développeur](#)
- [Jeux](#)
- [Internet des objets \(IoT\)](#)
- [Machine learning](#)
- [Management et gouvernance](#)
- [Services multimédias](#)
- [Migration et transfert](#)
- [Mise en réseau et diffusion de contenu](#)
- [Sécurité, identité et conformité](#)
- [Robotique](#)
- [Stockage](#)

Analytique

- Amazon AppFlow
- AWS Clean Rooms
- AWS Data Exchange
- Amazon DataZone
- Amazon Managed Streaming for Apache Kafka (Amazon MSK)
- Amazon Timestream pour LiveAnalytics

Intégration d'applications

- AWS AppFabric
- Amazon Simple Workflow Service

Applications métier

- Amazon WorkDocs
- Amazon WorkMail

Calcul

- AWS App Runner
- AWS Copilot
- AWS Wavelength

Gestion des coûts

- AWS Application Cost Profiler
- Amazon DevPay

Émancipation du client

- AWS Activate
- AWS IQ
- Amazon Managed Services (AMS)

Gestion financière du cloud

- AWS Billing Conductor

Base de données

- Amazon Keyspaces (for Apache Cassandra)

- Amazon MemoryDB for Redis OSS
- AWS AppConfig

Outils pour développeur

- AWS Application Composer
- AWS CodeArtifact
- AWS CodeDeploy
- Amazon CodeGuru
- AWS CloudShell
- AWS Device Farm

Jeux

- Amazon GameLift
- Amazon Lumberyard

Internet des objets (IoT)

- AWS IoT Device Defender
- AWS IoT Greengrass
- Amazon Monitron

Machine learning

- Amazon Fraud Detector
- Amazon Lookout for Metrics
- Amazon Mechanical Turk
- AWS Panorama
- Amazon Personalize

Management et gouvernance

- AWS Chatbot
- Amazon Data Lifecycle Manager
- Amazon Elastic Transcoder
- AWS Launch Wizard

Services multimédias

- AWS Elemental Appliances and Software
- AWS Elemental MediaConnect
- AWS Elemental MediaConvert
- AWS Elemental MediaLive
- AWS Elemental MediaPackage
- AWS Elemental MediaStore
- AWS Elemental MediaTailor
- Amazon Interactive Video Service (Amazon IVS)

Migration et transfert

- AWS Migration Hub Refactor Spaces
- AWS Transfer Family

Mise en réseau et diffusion de contenu

- AWS Cloud Map
- AWS Network Access Analyzer
- AWS Ground Station
- Amazon VPC Lattice

Sécurité, identité et conformité

- Amazon Cloud Directory

- AWS Network Firewall

Robotique

- AWS RoboMaker

Stockage

- Amazon FSx for Lustre

AWS Certified AI Practitioner (AIF-C01)

L'examen AWS Certified AI Practitioner (AIF-C01) est destiné aux personnes qui peuvent démontrer efficacement leur connaissance globale des technologies d'IA/ML et d'IA générative ainsi que des services et outils AWS associés, indépendamment d'un poste spécifique.

Rubriques

- [Introduction](#)
- [Description du candidat cible](#)
- [Contenu de l'examen](#)
- [Domaines de contenu](#)
- [Services AWS pour l'examen](#)
- [Domaine 1 : Principes fondamentaux de l'IA et du ML \(20 % du contenu noté\)](#)
- [Domaine 2 : Principes fondamentaux de l'IA générative \(24 % du contenu noté\)](#)
- [Domaine 3 : Applications des modèles de fondation \(28 % du contenu noté\)](#)
- [Domaine 4 : Directives pour une IA responsable \(14 % du contenu noté\)](#)
- [Domaine 5 : Sécurité, conformité et gouvernance pour les solutions d'IA \(14 % du contenu noté\)](#)
- [Services AWS concernés](#)
- [Services AWS non concernés](#)

Introduction

The AWS Certified AI Practitioner (AIF-C01) exam is intended for individuals who can effectively demonstrate overall knowledge of AI/ML, generative AI technologies, and associated AWS services and tools, independent of a specific job role.

The exam also validates a candidate's ability to complete the following tasks:

- Understand AI, ML, and generative AI concepts, methods, and strategies in general and on AWS.
- Understand the appropriate use of AI/ML and generative AI technologies to ask relevant questions within the candidate's organization.
- Determine the correct types of AI/ML technologies to apply to specific use cases.
- Use AI, ML, and generative AI technologies responsibly.

Description du candidat cible

Le candidat cible doit avoir jusqu'à 6 mois d'expérience des technologies d'IA/ML sur AWS. Le candidat cible utilise des solutions d'IA/ML sur AWS, mais ne les développe pas nécessairement.

Le candidat cible doit avoir les connaissances AWS suivantes :

- Connaissance des principaux services AWS (par exemple, Amazon EC2, Amazon S3, AWS Lambda et Amazon SageMaker) et des cas d'utilisation de ces principaux services AWS
- Connaissance du modèle de responsabilité partagée AWS en matière de sécurité et de conformité dans le cloud AWS
- Connaissance du service Gestion des identités et des accès AWS (AWS IAM) pour sécuriser et contrôler l'accès aux ressources AWS
- Connaissance de l'infrastructure globale AWS, y compris des concepts de régions AWS, de zones de disponibilité (AZ) et d'emplacements périphériques
- Connaissance des modèles de tarification des services AWS

La liste suivante contient les tâches que le candidat cible n'est pas censé être en mesure d'effectuer. Cette liste n'est pas exhaustive. Ces tâches sortent du champ d'application de l'examen :

- Développement ou codage de modèles ou d'algorithmes d'IA/ML
- Mise en œuvre de techniques d'ingénierie des données ou d'ingénierie des fonctionnalités

- Exécution du réglage des hyperparamètres ou de l'optimisation du modèle
- Création et déploiement de pipelines ou d'infrastructure d'IA/ML
- Réalisation d'analyses mathématiques ou statistiques de modèles d'IA/ML
- Mise en œuvre de protocoles de sécurité ou de conformité pour les systèmes d'IA/ML
- Élaboration et mise en œuvre de cadres et de politiques de gouvernance pour les solutions d'IA/ML

Contenu de l'examen

L'examen contient un ou plusieurs des types de questions suivants :

- Choix multiple : comprend une réponse correcte et trois réponses incorrectes (réponses piège).
- Réponse multiple : comprend au moins deux réponses correctes sur cinq options de réponse ou plus. Vous devez sélectionner toutes les bonnes réponses pour valider la question.
- Classement : comprend une liste de 3 à 5 réponses pour effectuer une tâche spécifiée. Vous devez sélectionner les bonnes réponses et les placer dans le bon ordre pour valider la question.
- Mise en correspondance : comprend une liste de réponses à mettre en correspondance avec une liste de 3 à 7 invites. Vous devez associer toutes les paires correctement pour valider la question.
- Étude de cas : comprend un scénario et au moins deux questions sur ce scénario. Le scénario est le même pour chaque question de l'étude de cas. Chaque question de l'étude de cas est évaluée séparément. Vous validez chaque question à laquelle vous répondez correctement dans l'étude de cas.

Les questions laissées sans réponse sont notées comme incorrectes ; aucune pénalité n'est appliquée lorsque le candidat devine une réponse. L'examen comporte 50 questions qui ont une incidence sur votre score.

L'examen comporte 15 questions non notées qui n'ont pas d'incidence sur votre score. AWS recueille des informations sur les performances dans le cadre de ces questions non notées. L'objectif est d'évaluer ces questions en vue d'une utilisation ultérieure en tant que questions notées. Ces questions non notées ne sont pas identifiées comme telles dans l'examen.

L'examen AWS Certified AI Practitioner (AIF-C01) est sanctionné par une réussite ou un échec. L'examen est évalué selon une norme minimale établie par les professionnels d'AWS. Ceux-ci observent les bonnes pratiques et directives en matière de certification.

Vos résultats à l'examen sont présentés sous la forme d'un score gradué de 100 à 1 000. Le score minimal pour réussir est de 700. Votre score indique vos performances lors de l'examen dans son ensemble et si vous avez réussi l'examen. Les modèles de notation graduée permettent de mettre en correspondance des scores de différents formulaires d'examen qui peuvent présenter des niveaux de difficulté légèrement différents.

Votre compte rendu de score peut contenir un tableau des classifications de vos performances au niveau de chaque section. L'examen utilise un modèle de notation compensatoire, ce qui signifie que vous n'avez pas besoin d'obtenir une note minimale dans chaque section. Vous devez seulement réussir l'examen dans son ensemble.

Chaque section de l'examen présente une pondération spécifique. Certaines sections comportent donc plus de questions que d'autres. Le tableau des classifications contient des informations générales qui mettent en évidence vos points forts et vos points faibles. Les commentaires au niveau des sections doivent être interprétés avec précaution.

Domaines de contenu

Ce guide de l'examen inclut les pondérations, les domaines du contenu ainsi que les énoncés de tâche de l'examen. Ce guide ne fournit pas de liste exhaustive du contenu de l'examen. Cependant, chaque énoncé de tâche est accompagné d'un contexte supplémentaire pour vous aider à vous préparer à l'examen.

L'examen comporte les domaines de contenu et les pondérations suivants :

- Domaine 1 : Principes fondamentaux de l'IA et du ML (20 % du contenu noté)
- Domaine 2 : Principes fondamentaux de l'IA générative (24 % du contenu noté)
- Domaine 3 : Applications des modèles de fondation (28 % du contenu noté)
- Domaine 4 : Directives pour une IA responsable (14 % du contenu noté)
- Domaine 5 : Sécurité, conformité et gouvernance pour les solutions d'IA (14 % du contenu évalué)

Pour des informations détaillées sur chaque domaine, consultez les sections suivantes :

- [the section called “Domaine 1 : Principes fondamentaux de l'IA et du ML \(20 % du contenu noté\)”](#)
- [the section called “Domaine 2 : Principes fondamentaux de l'IA générative \(24 % du contenu noté\)”](#)
- [the section called “Domaine 3 : Applications des modèles de fondation \(28 % du contenu noté\)”](#)
- [the section called “Domaine 4 : Directives pour une IA responsable \(14 % du contenu noté\)”](#)

- [the section called “Domaine 5 : Sécurité, conformité et gouvernance pour les solutions d'IA \(14 % du contenu noté\)”](#)

Services AWS pour l'examen

L'examen AWS Certified AI Practitioner couvre des services AWS spécifiques qui sont pertinents pour les charges de travail d'IA et de ML. Comprendre quels services sont dans le champ d'application et lesquels sont hors du champ d'application peut vous aider à concentrer vos efforts de préparation.

Pour des informations détaillées sur les services AWS couverts dans l'examen, consultez les sections suivantes :

- [Services AWS dans le champ d'application](#)
- [Services AWS hors du champ d'application](#)

Domaine 1 : Principes fondamentaux de l'IA et du ML (20 % du contenu noté)

Le Domaine 1 couvre les principes fondamentaux de l'IA et du ML et représente 20 % du contenu noté de l'examen.

Rubriques

- [Énoncé de tâche 1.1 : Expliquer les concepts et terminologies de base de l'IA.](#)
- [Énoncé de tâche 1.2 : Identifier des cas d'utilisation pratiques de l'IA.](#)
- [Énoncé de tâche 1.3 : Décrire le cycle de vie du développement du machine learning.](#)

Énoncé de tâche 1.1 : Expliquer les concepts et terminologies de base de l'IA.

Objectifs :

- Définir les termes de base de l'IA (par exemple, IA, ML, deep learning, réseaux neuronaux, vision par ordinateur, traitement du langage naturel [NLP], modèle, algorithme, entraînement et inférence, biais, équité, ajustement, grand modèle de langage [LLM]).
- Décrire les similitudes et les différences entre l'IA, le ML et le deep learning.
- Décrire les différents types d'inférence (par exemple, par lots, en temps réel).

- Décrire les différents types de données dans les modèles d'IA (par exemple, étiquetées et non étiquetées, tabulaires, séries chronologiques, images, textes, structurées et non structurées).
- Décrire l'apprentissage supervisé, l'apprentissage non supervisé et l'apprentissage par renforcement.

Énoncé de tâche 1.2 : Identifier des cas d'utilisation pratiques de l'IA.

Objectifs :

- Identifier les applications dans lesquelles l'IA ou le ML peut apporter de la valeur (par exemple, aide à la prise de décision humaine, capacité de mise à l'échelle des solutions, automatisation).
- Déterminer quand les solutions d'IA/ML ne sont pas appropriées (par exemple, analyses coûts-avantages, situations dans lesquelles un résultat spécifique est nécessaire au lieu d'une prédiction).
- Sélectionner les techniques de ML appropriées pour des cas d'utilisation spécifiques (par exemple, régression, classification, mise en cluster).
- Identifier des exemples d'applications d'IA concrètes (par exemple, vision par ordinateur, NLP, reconnaissance vocale, systèmes de recommandation, détection de fraude, prévisions).
- Expliquer les fonctionnalités des services d'IA/ML gérés par AWS (par exemple, SageMaker, Amazon Transcribe, Amazon Translate, Amazon Comprehend, Amazon Lex, Amazon Polly).

Énoncé de tâche 1.3 : Décrire le cycle de vie du développement du machine learning.

Objectifs :

- Décrire les composants d'un pipeline de ML (par exemple, collecte de données, analyse exploratoire des données [EDA], prétraitement des données, ingénierie des fonctionnalités, entraînement de modèle, réglage des hyperparamètres, évaluation, déploiement, surveillance).
- Comprendre les sources des modèles de ML (par exemple, modèles open source pré-entraînés, modèles d'entraînement personnalisés).
- Décrire les méthodes d'utilisation d'un modèle en production (par exemple, service d'API géré, API auto-hébergée).
- Identifier les services et fonctionnalités AWS pertinents pour chaque étape d'un pipeline de ML (par exemple, SageMaker, Amazon SageMaker Data Wrangler, Amazon SageMaker Feature Store, Amazon SageMaker Model Monitor).

- Comprendre les concepts fondamentaux des opérations de ML (MLOps) (par exemple, expérimentation, processus répétables, systèmes pouvant être mis à l'échelle, gestion de la dette technique, préparation à la production, surveillance des modèles, réentraînement des modèles).
- Comprendre les métriques de performance des modèles (par exemple, précision, aire sous la courbe ROC [AUC], score F1) et les métriques opérationnelles (par exemple, coût par utilisateur, coûts de développement, commentaires des clients, retour sur investissement [ROI]) pour évaluer les modèles de ML.

Domaine 2 : Principes fondamentaux de l'IA générative (24 % du contenu noté)

Le Domaine 2 couvre les principes fondamentaux de l'IA générative et représente 24 % du contenu noté de l'examen.

Rubriques

- [Task Statement 2.1: Explain the basic concepts of generative AI](#)
- [Task Statement 2.2: Understand the capabilities and limitations of generative AI for solving business problems](#)
- [Task Statement 2.3: Describe AWS infrastructure and technologies for building generative AI applications](#)

Task Statement 2.1: Explain the basic concepts of generative AI

Objectives:

- Understand foundational generative AI concepts (for example, tokens, chunking, embeddings, vectors, prompt engineering, transformer-based LLMs, foundation models, multi-modal models, diffusion models).
- Identify potential use cases for generative AI models (for example, image, video, and audio generation; summarization; chatbots; translation; code generation; customer service agents; search; recommendation engines).
- Describe the foundation model lifecycle (for example, data selection, model selection, pre-training, fine-tuning, evaluation, deployment, feedback).

Task Statement 2.2: Understand the capabilities and limitations of generative AI for solving business problems

Objectives:

- Describe the advantages of generative AI (for example, adaptability, responsiveness, simplicity).
- Identify disadvantages of generative AI solutions (for example, hallucinations, interpretability, inaccuracy, nondeterminism).
- Understand various factors to select appropriate generative AI models (for example, model types, performance requirements, capabilities, constraints, compliance).
- Determine business value and metrics for generative AI applications (for example, cross-domain performance, efficiency, conversion rate, average revenue per user, accuracy, customer lifetime value).

Task Statement 2.3: Describe AWS infrastructure and technologies for building generative AI applications

Objectives:

- Identify AWS services and features to develop generative AI applications (for example, Amazon SageMaker JumpStart; Amazon Bedrock; PartyRock, an Amazon Bedrock Playground; Amazon Q).
- Describe the advantages of using AWS generative AI services to build applications (for example, accessibility, lower barrier to entry, efficiency, cost-effectiveness, speed to market, ability to meet business objectives).
- Understand the benefits of AWS infrastructure for generative AI applications (for example, security, compliance, responsibility, safety).
- Understand cost tradeoffs of AWS generative AI services (for example, responsiveness, availability, redundancy, performance, regional coverage, token-based pricing, provision throughput, custom models).

Domaine 3 : Applications des modèles de fondation (28 % du contenu noté)

Le Domaine 3 couvre les applications des modèles de fondation et représente 28 % du contenu noté de l'examen.

Rubriques

- [Task Statement 3.1: Describe design considerations for applications that use foundation models](#)
- [Task Statement 3.2: Choose effective prompt engineering techniques](#)
- [Task Statement 3.3: Describe the training and fine-tuning process for foundation models](#)
- [Task Statement 3.4: Describe methods to evaluate foundation model performance](#)

Task Statement 3.1: Describe design considerations for applications that use foundation models

Objectives:

- Identify selection criteria to choose pre-trained models (for example, cost, modality, latency, multi-lingual, model size, model complexity, customization, input/output length).
- Understand the effect of inference parameters on model responses (for example, temperature, input/output length).
- Define Retrieval Augmented Generation (RAG) and describe its business applications (for example, Amazon Bedrock, knowledge base).
- Identify AWS services that help store embeddings within vector databases (for example, Amazon OpenSearch Service, Amazon Aurora, Amazon Neptune, Amazon DocumentDB [with MongoDB compatibility], Amazon RDS for PostgreSQL).
- Explain the cost tradeoffs of various approaches to foundation model customization (for example, pre-training, fine-tuning, in-context learning, RAG).
- Understand the role of agents in multi-step tasks (for example, Agents for Amazon Bedrock).

Task Statement 3.2: Choose effective prompt engineering techniques

Objectives:

- Describe the concepts and constructs of prompt engineering (for example, context, instruction, negative prompts, model latent space).
- Understand techniques for prompt engineering (for example, chain-of-thought, zero-shot, single-shot, few-shot, prompt templates).
- Understand the benefits and best practices for prompt engineering (for example, response quality improvement, experimentation, guardrails, discovery, specificity and concision, using multiple comments).

- Define potential risks and limitations of prompt engineering (for example, exposure, poisoning, hijacking, jailbreaking).

Task Statement 3.3: Describe the training and fine-tuning process for foundation models

Objectives:

- Describe the key elements of training a foundation model (for example, pre-training, fine-tuning, continuous pre-training).
- Define methods for fine-tuning a foundation model (for example, instruction tuning, adapting models for specific domains, transfer learning, continuous pre-training).
- Describe how to prepare data to fine-tune a foundation model (for example, data curation, governance, size, labeling, representativeness, reinforcement learning from human feedback [RLHF]).

Task Statement 3.4: Describe methods to evaluate foundation model performance

Objectives:

- Understand approaches to evaluate foundation model performance (for example, human evaluation, benchmark datasets).
- Identify relevant metrics to assess foundation model performance (for example, Recall-Oriented Understudy for Gisting Evaluation [ROUGE], Bilingual Evaluation Understudy [BLEU], BERTScore).
- Determine whether a foundation model effectively meets business objectives (for example, productivity, user engagement, task engineering).

Domaine 4 : Directives pour une IA responsable (14 % du contenu noté)

Le Domaine 4 couvre les directives pour une IA responsable et représente 14 % du contenu noté de l'examen.

Rubriques

- [Task Statement 4.1: Explain the development of AI systems that are responsible](#)
- [Task Statement 4.2: Recognize the importance of transparent and explainable models](#)

Task Statement 4.1: Explain the development of AI systems that are responsible

Objectives:

- Identify features of responsible AI (for example, bias, fairness, inclusivity, robustness, safety, veracity).
- Understand how to use tools to identify features of responsible AI (for example, Guardrails for Amazon Bedrock).
- Understand responsible practices to select a model (for example, environmental considerations, sustainability).
- Identify legal risks of working with generative AI (for example, intellectual property infringement claims, biased model outputs, loss of customer trust, end user risk, hallucinations).
- Identify characteristics of datasets (for example, inclusivity, diversity, curated data sources, balanced datasets).
- Understand effects of bias and variance (for example, effects on demographic groups, inaccuracy, overfitting, underfitting).
- Describe tools to detect and monitor bias, trustworthiness, and truthfulness (for example, analyzing label quality, human audits, subgroup analysis, Amazon SageMaker Clarify, SageMaker Model Monitor, Amazon Augmented AI [Amazon A2I]).

Task Statement 4.2: Recognize the importance of transparent and explainable models

Objectives:

- Understand the differences between models that are transparent and explainable and models that are not transparent and explainable.
- Understand the tools to identify transparent and explainable models (for example, Amazon SageMaker Model Cards, open source models, data, licensing).
- Identify tradeoffs between model safety and transparency (for example, measure interpretability and performance).
- Understand principles of human-centered design for explainable AI.

Domaine 5 : Sécurité, conformité et gouvernance pour les solutions d'IA (14 % du contenu noté)

Le Domaine 5 couvre la sécurité, la conformité et la gouvernance pour les solutions d'IA et représente 14 % du contenu noté de l'examen.

Rubriques

- [Task Statement 5.1: Explain methods to secure AI systems](#)
- [Task Statement 5.2: Recognize governance and compliance regulations for AI systems](#)

Task Statement 5.1: Explain methods to secure AI systems

Objectives:

- Identify AWS services and features to secure AI systems (for example, IAM roles, policies, and permissions; encryption; Amazon Macie; AWS PrivateLink; AWS shared responsibility model).
- Understand the concept of source citation and documenting data origins (for example, data lineage, data cataloging, SageMaker Model Cards).
- Describe best practices for secure data engineering (for example, assessing data quality, implementing privacy-enhancing technologies, data access control, data integrity).
- Understand security and privacy considerations for AI systems (for example, application security, threat detection, vulnerability management, infrastructure protection, prompt injection, encryption at rest and in transit).

Task Statement 5.2: Recognize governance and compliance regulations for AI systems

Objectives:

- Identify regulatory compliance standards for AI systems (for example, International Organization for Standardization [ISO], System and Organization Controls [SOC], algorithm accountability laws).
- Identify AWS services and features to assist with governance and regulation compliance (for example, AWS Config, Amazon Inspector, AWS Audit Manager, AWS Artifact, AWS CloudTrail, AWS Trusted Advisor).
- Describe data governance strategies (for example, data lifecycles, logging, residency, monitoring, observation, retention).

- Describe processes to follow governance protocols (for example, policies, review cadence, review strategies, governance frameworks such as the Generative AI Security Scoping Matrix, transparency standards, team training requirements).

Services AWS concernés

Rubriques

- [Analytique](#)
- [Gestion financière du cloud](#)
- [Calcul](#)
- [Conteneurs](#)
- [Base de données](#)
- [Machine learning](#)
- [Management et gouvernance](#)
- [Mise en réseau et diffusion de contenu](#)
- [Sécurité, identité et conformité](#)
- [Stockage](#)

La liste suivante contient les services et fonctionnalités AWS qui font partie du champ d'application de l'examen AWS Certified AI Practitioner (AIF-C01). Cette liste n'est pas exhaustive et peut faire l'objet de modifications. Les offres AWS apparaissent dans des catégories correspondant à leurs fonctions principales.

Analytique

- AWS Data Exchange
- Amazon EMR
- AWS Glue
- AWS Glue DataBrew
- AWS Lake Formation
- Amazon OpenSearch Service
- Amazon QuickSight
- Amazon Redshift

Gestion financière du cloud

- AWS Budgets
- AWS Cost Explorer

Calcul

- Amazon EC2

Conteneurs

- Amazon Elastic Container Service (Amazon ECS)
- Amazon Elastic Kubernetes Service (Amazon EKS)

Base de données

- Amazon DocumentDB (compatible avec MongoDB)
- Amazon DynamoDB
- Amazon ElastiCache
- Amazon MemoryDB
- Amazon Neptune
- Amazon RDS

Machine learning

- Amazon Augmented AI (Amazon A2I)
- Amazon Bedrock
- Amazon Comprehend
- Amazon Fraud Detector
- Amazon Kendra
- Amazon Lex
- Amazon Personalize
- Amazon Polly

- Amazon Q
- Amazon Rekognition
- Amazon SageMaker
- Amazon Textract
- Amazon Transcribe
- Amazon Translate

Management et gouvernance

- AWS CloudTrail
- Amazon CloudWatch
- AWS Config
- AWS Trusted Advisor
- Outil AWS Well-Architected

Mise en réseau et diffusion de contenu

- Amazon CloudFront
- Amazon VPC

Sécurité, identité et conformité

- AWS Artifact
- AWS Audit Manager
- Gestion des identités et des accès AWS (IAM)
- Amazon Inspector
- AWS Key Management Service (AWS KMS)
- Amazon Macie
- AWS Secrets Manager

Stockage

- Amazon S3

- Amazon S3 Glacier

Services AWS non concernés

Rubriques

- [Analytique](#)
- [Intégration d'applications](#)
- [Applications métier](#)
- [Gestion financière du cloud](#)
- [Calcul](#)
- [Conteneurs](#)
- [Émancipation du client](#)
- [Base de données](#)
- [Outils pour développeur](#)
- [Informatique pour l'utilisateur final](#)
- [Applications web et mobiles front-end](#)
- [Internet des objets \(IoT\)](#)
- [Machine learning](#)
- [Management et gouvernance](#)
- [Médias](#)
- [Migration et transfert](#)
- [Mise en réseau et diffusion de contenu](#)
- [Sécurité, identité et conformité](#)
- [Stockage](#)

La liste suivante contient les services et fonctionnalités AWS qui sont hors du champ d'application de l'examen AWS Certified AI Practitioner (AIF-C01). Cette liste n'est pas exhaustive et peut faire l'objet de modifications. Les offres AWS qui ne sont aucunement liées aux rôles cibles de l'examen sont exclues de cette liste.

Analytique

- AWS Clean Rooms

- Amazon CloudSearch
- Amazon FinSpace
- Amazon Managed Streaming for Apache Kafka (Amazon MSK)

Intégration d'applications

- Amazon AppFlow
- Amazon MQ
- Amazon Simple Workflow Service (Amazon SWF)

Applications métier

- Amazon Chime
- Amazon Honeycode
- Amazon Pinpoint
- Amazon Simple Email Service (Amazon SES)
- Chaîne d'approvisionnement AWS
- AWS Wickr
- Amazon WorkDocs
- Amazon WorkMail

Gestion financière du cloud

- AWS Application Cost Profiler
- AWS Billing Conductor
- AWS Marketplace

Calcul

- AWS App Runner
- AWS Elastic Beanstalk
- EC2 Image Builder
- Amazon Lightsail

Conteneurs

- Service Red Hat OpenShift sur AWS (ROSA)

Émancipation du client

- AWS IQ
- AWS Managed Services (AMS)
- AWS re:Post Private
- AWS Support

Base de données

- Amazon Keyspaces (pour Apache Cassandra)
- Amazon Quantum Ledger Database (Amazon QLDB)
- Amazon Timestream

Outils pour développeur

- AWS AppConfig
- AWS Application Composer
- AWS CloudShell
- Amazon CodeCatalyst
- AWS CodeStar
- AWS Fault Injection Service
- AWS X-Ray

Informatique pour l'utilisateur final

- Amazon AppStream 2.0
- Amazon WorkSpaces
- Client léger Amazon WorkSpaces
- Amazon WorkSpaces Web

Applications web et mobiles front-end

- AWS Amplify
- AWS AppSync
- AWS Device Farm
- Amazon Location Service

Internet des objets (IoT)

- AWS IoT Analytics
- AWS IoT Core
- AWS IoT Device Defender
- AWS IoT Device Management
- AWS IoT Events
- AWS IoT FleetWise
- FreeRTOS
- AWS IoT Greengrass
- AWS IoT 1-Click
- AWS IoT RoboRunner
- AWS IoT SiteWise
- AWS IoT TwinMaker

Machine learning

- AWS DeepComposer
- AWS HealthImaging
- AWS HealthOmics
- Amazon Monitron
- AWS Panorama

Management et gouvernance

- AWS Control Tower

- Tableau de bord AWS Health
- AWS Launch Wizard
- AWS License Manager
- Amazon Managed Grafana
- Amazon Managed Service for Prometheus
- AWS OpsWorks
- AWS Organizations
- AWS Proton
- AWS Resilience Hub
- Explorateur de ressources AWS
- AWS Resource Groups
- AWS Systems Manager Incident Manager
- AWS Service Catalog
- Service Quotas
- Créateur de réseau de télécommunications AWS
- Notifications des utilisateurs AWS

Médias

- Amazon Elastic Transcoder
- AWS Elemental MediaConnect
- AWS Elemental MediaConvert
- AWS Elemental MediaLive
- AWS Elemental MediaPackage
- AWS Elemental MediaStore
- AWS Elemental MediaTailor
- Amazon Interactive Video Service (Amazon IVS)
- Amazon Nimble Studio

Migration et transfert

- AWS Application Discovery Service

- AWS Application Migration Service
- AWS Database Migration Service (AWS DMS)
- AWS DataSync
- AWS Mainframe Modernization
- AWS Migration Hub
- Gamme AWS Snow
- AWS Transfer Family

Mise en réseau et diffusion de contenu

- AWS App Mesh
- AWS Cloud Map
- AWS Direct Connect
- AWS Global Accelerator
- Services privés 5G d'AWS
- Amazon Route 53
- Contrôleur de récupération d'application Amazon Route 53
- Gestionnaire d'adresses IP Amazon VPC (IPAM)

Sécurité, identité et conformité

- AWS Certificate Manager (ACM)
- AWS CloudHSM
- Amazon Cognito
- Amazon Detective
- AWS Directory Service
- AWS Firewall Manager
- Amazon GuardDuty
- AWS IAM Identity Center
- AWS Payment Cryptography
- Autorité de Certification Privée AWS
- AWS Resource Access Manager (AWS RAM)

- AWS Security Hub
- Amazon Security Lake
- AWS Shield
- AWS Signer
- Amazon Verified Permissions
- AWS WAF

Stockage

- AWS Backup
- AWS Elastic Disaster Recovery

Associate Exams

AWS Associate level certification exams validate technical skills and experience in implementing AWS solutions.

Rubriques

- [AWS Certified Developer - Associate \(DVA-C02\)](#)
- [AWS Certified SysOps Administrator - Associate \(SOA-C02\)](#)
- [AWS Certified Solutions Architect - Associate \(SAA-C03\)](#)
- [AWS Certified Machine Learning Engineer - Associate \(MLA-C01\)](#)
- [AWS Certified Data Engineer - Associate \(DEA-C01\)](#)

AWS Certified Developer - Associate (DVA-C02)

The AWS Certified Developer - Associate (DVA-C02) exam is intended for individuals who perform a developer role. The exam validates a candidate's ability to demonstrate proficiency in developing, testing, deploying, and debugging AWS Cloud-based applications.

Rubriques

- [Introduction](#)
- [Target Candidate Description](#)
- [Exam Content](#)
- [Content Domains](#)
- [AWS Services for the Exam](#)
- [Domain 1: Development with Services \(32% of the exam content\)](#)
- [Domain 2: Security \(26% of the exam content\)](#)
- [Domain 3: Deployment \(24% of the exam content\)](#)
- [Domain 4: Troubleshooting and Optimization \(18% of the exam content\)](#)
- [In-Scope AWS Services](#)
- [Out-of-Scope AWS Services](#)

Introduction

The AWS Certified Developer - Associate (DVA-C02) exam is intended for individuals who perform a developer role. The exam validates a candidate's ability to demonstrate proficiency in developing, testing, deploying, and debugging AWS Cloud-based applications.

The exam also validates a candidate's ability to complete the following tasks:

- Develop and optimize applications on AWS.
- Package and deploy by using continuous integration and continuous delivery (CI/CD) workflows.
- Secure application code and data.
- Identify and resolve application issues.

Target Candidate Description

The target candidate should have 1 or more years of hands-on experience in developing and maintaining applications by using AWS services.

The target candidate should have the following general IT knowledge:

- Proficiency in at least one high-level programming language
- Understanding of application lifecycle management
- Basic understanding of cloud-focused applications to write code
- Ability to develop functional applications
- Experience in using development tools

The target candidate should be able to complete the following tasks:

- Develop and secure applications by using AWS service APIs, the AWS Command Line Interface (AWS CLI), and SDKs.
- Use a CI/CD pipeline to deploy applications on AWS.

The following list contains job tasks that the target candidate is not expected to be able to perform. This list is non-exhaustive. These tasks are out of scope for the exam:

- Design architectures (for example, distributed systems, microservices, database schemas and modeling).
- Design and create CI/CD pipelines.
- Administer IAM users and groups.
- Administer servers and operating systems.
- Design AWS networking infrastructure (for example, Amazon Virtual Private Cloud [Amazon VPC], AWS Direct Connect).

Exam Content

There are two types of questions on the exam:

- Multiple choice: Has one correct response and three incorrect responses (distractors)
- Multiple response: Has two or more correct responses out of five or more response options

Select one or more responses that best complete the statement or answer the question. Distractors, or incorrect answers, are response options that a candidate with incomplete knowledge or skill might choose. Distractors are generally plausible responses that match the content area.

Unanswered questions are scored as incorrect; there is no penalty for guessing. The exam includes 50 questions that affect your score.

The exam includes 15 unscored questions that do not affect your score. AWS collects information about performance on these unscored questions to evaluate these questions for future use as scored questions. These unscored questions are not identified on the exam.

The AWS Certified Developer - Associate (DVA-C02) exam has a pass or fail designation. The exam is scored against a minimum standard established by AWS professionals who follow certification industry best practices and guidelines.

Your results for the exam are reported as a scaled score of 100–1,000. The minimum passing score is 720. Your score shows how you performed on the exam as a whole and whether you passed. Scaled scoring models help equate scores across multiple exam forms that might have slightly different difficulty levels.

Content Domains

For detailed information about each domain, see the following sections:

- [the section called “Domain 1: Development with Services \(32% of the exam content\)”](#)
- [the section called “Domain 2: Security \(26% of the exam content\)”](#)
- [the section called “Domain 3: Deployment \(24% of the exam content\)”](#)
- [the section called “Domain 4: Troubleshooting and Optimization \(18% of the exam content\)”](#)

AWS Services for the Exam

The AWS Certified Developer - Associate exam covers specific AWS services that are relevant to developing applications on AWS. Understanding which services are in scope and which are out of scope can help you focus your preparation efforts.

For detailed information about the AWS services covered in the exam, see the following sections:

- [In-Scope AWS Services](#)
- [Out-of-Scope AWS Services](#)

Domain 1: Development with Services (32% of the exam content)

This domain accounts for 32% of the exam content.

Rubriques

- [Task 1: Develop code for applications hosted on](#)
- [Task 2: Develop code for Lambda](#)
- [Task 3: Use data stores in application development](#)

Task 1: Develop code for applications hosted on

Knowledge of:

- Architectural patterns (for example, event-driven, microservices, monolithic, choreography, orchestration, fanout)
- Idempotency
- Differences between stateful and stateless concepts
- Differences between tightly coupled and loosely coupled components

- Fault-tolerant design patterns (for example, retries with exponential backoff and jitter, dead-letter queues)
- Differences between synchronous and asynchronous patterns

Skills in:

- Creating fault-tolerant and resilient applications in a programming language (for example, Java, C#, Python, JavaScript, TypeScript, Go)
- Creating, extending, and maintaining APIs (for example, response/request transformations, enforcing validation rules, overriding status codes)
- Writing and running unit tests in development environments (for example, using Serverless Application Model [SAM])
- Writing code to use messaging services
- Writing code that interacts with services by using APIs and SDKs
- Handling data streaming by using services

Task 2: Develop code for Lambda

Knowledge of:

- Event source mapping
- Stateless applications
- Unit testing
- Event-driven architecture
- Scalability
- The access of private resources in VPCs from Lambda code

Skills in:

- Configuring Lambda functions by defining environment variables and parameters (for example, memory, concurrency, timeout, runtime, handler, layers, extensions, triggers, destinations)
- Handling the event lifecycle and errors by using code (for example, Lambda Destinations, dead-letter queues)
- Writing and running test code by using services and tools

- Integrating Lambda functions with services
- Tuning Lambda functions for optimal performance

Task 3: Use data stores in application development

Knowledge of:

- Relational and non-relational databases
- Create, read, update, and delete (CRUD) operations
- High-cardinality partition keys for balanced partition access
- Cloud storage options (for example, file, object, databases)
- Database consistency models (for example, strongly consistent, eventually consistent)
- Differences between query and scan operations
- Amazon DynamoDB keys and indexing
- Caching strategies (for example, write-through, read-through, lazy loading, TTL)
- Amazon Simple Storage Service (Amazon S3) tiers and lifecycle management
- Differences between ephemeral and persistent data storage patterns

Skills in:

- Serializing and deserializing data to provide persistence to a data store
- Using, managing, and maintaining data stores
- Managing data lifecycles
- Using data caching services

Domain 2: Security (26% of the exam content)

This domain accounts for 26% of the exam content.

Rubriques

- [Task 1: Implement authentication and/or authorization for applications and services](#)
- [Task 2: Implement encryption by using services](#)

- [Task 3: Manage sensitive data in application code](#)

Task 1: Implement authentication and/or authorization for applications and services

Knowledge of:

- Identity federation (for example, Security Assertion Markup Language [SAML], OpenID Connect [OIDC], Amazon Cognito)
- Bearer tokens (for example, JSON Web Token [JWT], OAuth, Security Token Service [STS])
- The comparison of user pools and identity pools in Amazon Cognito
- Resource-based policies, service policies, and principal policies
- Role-based access control (RBAC)
- Application authorization that uses ACLs
- The principle of least privilege
- Differences between managed policies and customer-managed policies
- Identity and access management

Skills in:

- Using an identity provider to implement federated access (for example, Amazon Cognito, Identity and Access Management [IAM])
- Securing applications by using bearer tokens
- Configuring programmatic access to
- Making authenticated calls to services
- Assuming an IAM role
- Defining permissions for principals

Task 2: Implement encryption by using services

Knowledge of:

- Encryption at rest and in transit
- Certificate management (for example, Private Certificate Authority)

- Key protection (for example, key rotation)
- Differences between client-side encryption and server-side encryption
- Differences between managed and customer managed Key Management Service (KMS) keys

Skills in:

- Using encryption keys to encrypt or decrypt data
- Generating certificates and SSH keys for development purposes
- Using encryption across account boundaries
- Enabling and disabling key rotation

Task 3: Manage sensitive data in application code

Knowledge of:

- Data classification (for example, personally identifiable information [PII], protected health information [PHI])
- Environment variables
- Secrets management (for example, Secrets Manager, Systems Manager Parameter Store)
- Secure credential handling

Skills in:

- Encrypting environment variables that contain sensitive data
- Using secret management services to secure sensitive data
- Sanitizing sensitive data

Domain 3: Deployment (24% of the exam content)

This domain accounts for 24% of the exam content.

Rubriques

- [Task 1: Prepare application artifacts to be deployed to](#)
- [Task 2: Test applications in development environments](#)

- [Task 3: Automate deployment testing](#)
- [Task 4: Deploy code by using CI/CD services](#)

Task 1: Prepare application artifacts to be deployed to

Knowledge of:

- Ways to access application configuration data (for example, AppConfig, Secrets Manager, Parameter Store)
- Lambda deployment packaging, layers, and configuration options
- Git-based version control tools (for example, Git)
- Container images

Skills in:

- Managing the dependencies of the code module (for example, environment variables, configuration files, container images) within the package
- Organizing files and a directory structure for application deployment
- Using code repositories in deployment environments
- Applying application requirements for resources (for example, memory, cores)

Task 2: Test applications in development environments

Knowledge of:

- Features in services that perform application deployment
- Integration testing that uses mock endpoints
- Lambda versions and aliases

Skills in:

- Testing deployed code by using services and tools
- Performing mock integration for APIs and resolving integration dependencies
- Testing applications by using development endpoints (for example, configuring stages in Amazon API Gateway)

- Deploying application stack updates to existing environments (for example, deploying an SAM template to a different staging environment)

Task 3: Automate deployment testing

Knowledge of:

- API Gateway stages
- Branches and actions in the continuous integration and continuous delivery (CI/CD) workflow
- Automated software testing (for example, unit testing, mock testing)

Skills in:

- Creating application test events (for example, JSON payloads for testing Lambda, API Gateway, SAM resources)
- Deploying API resources to various environments
- Creating application environments that use approved versions for integration testing (for example, Lambda aliases, container image tags, Amplify branches, Copilot environments)
- Implementing and deploying infrastructure as code (IaC) templates (for example, SAM templates, CloudFormation templates)
- Managing environments in individual services (for example, differentiating between development, test, and production in API Gateway)

Task 4: Deploy code by using CI/CD services

Knowledge of:

- Git-based version control tools (for example, Git)
- Manual and automated approvals in CodePipeline
- Access application configurations from AppConfig and Secrets Manager
- CI/CD workflows that use services
- Application deployment that uses services and tools (for example, CloudFormation, Cloud Development Kit [CDK], SAM, CodeArtifact, Copilot, Amplify, Lambda)
- Lambda deployment packaging options

- API Gateway stages and custom domains
- Deployment strategies (for example, canary, blue/green, rolling)

Skills in:

- Updating existing IaC templates (for example, SAM templates, CloudFormation templates)
- Managing application environments by using services
- Deploying an application version by using deployment strategies
- Committing code to a repository to invoke build, test, and deployment actions
- Using orchestrated workflows to deploy code to different environments
- Performing application rollbacks by using existing deployment strategies
- Using labels and branches for version and release management
- Using existing runtime configurations to create dynamic deployments (for example, using staging variables from API Gateway in Lambda functions)

Domain 4: Troubleshooting and Optimization (18% of the exam content)

This domain accounts for 18% of the exam content.

Rubriques

- [Task 1: Assist in a root cause analysis](#)
- [Task 2: Instrument code for observability](#)
- [Task 3: Optimize applications by using services and features](#)

Task 1: Assist in a root cause analysis

Knowledge of:

- Logging and monitoring systems
- Languages for log queries (for example, Amazon CloudWatch Logs Insights)
- Data visualizations
- Code analysis tools
- Common HTTP error codes

- Common exceptions generated by SDKs
- Service maps in X-Ray

Skills in:

- Debugging code to identify defects
- Interpreting application metrics, logs, and traces
- Querying logs to find relevant data
- Implementing custom metrics (for example, CloudWatch embedded metric format [EMF])
- Reviewing application health by using dashboards and insights
- Troubleshooting deployment failures by using service output logs

Task 2: Instrument code for observability

Knowledge of:

- Distributed tracing
- Differences between logging, monitoring, and observability
- Structured logging
- Application metrics (for example, custom, embedded, built-in)

Skills in:

- Implementing an effective logging strategy to record application behavior and state
- Implementing code that emits custom metrics
- Adding annotations for tracing services
- Implementing notification alerts for specific actions (for example, notifications about quota limits or deployment completions)
- Implementing tracing by using services and tools

Task 3: Optimize applications by using services and features

Knowledge of:

- Caching
- Concurrency
- Messaging services (for example, Amazon Simple Queue Service [Amazon SQS], Amazon Simple Notification Service [Amazon SNS])

Skills in:

- Profiling application performance
- Determining minimum memory and compute power for an application
- Using subscription filter policies to optimize messaging
- Caching content based on request headers

In-Scope AWS Services

The following list contains AWS services and features that are in scope for the AWS Certified Developer - Associate (DVA-C02) exam. This list is non-exhaustive and is subject to change. AWS offerings appear in categories that align with the offerings' primary functions.

Rubriques

- [Analytics](#)
- [Application Integration](#)
- [Compute](#)
- [Containers](#)
- [Database](#)
- [Developer Tools](#)
- [Management and Governance](#)
- [Networking and Content Delivery](#)
- [Security, Identity, and Compliance](#)
- [Storage](#)

Analytics

- Amazon Athena

- Amazon Kinesis
- Amazon OpenSearch Service

Application Integration

- Amazon EventBridge
- Amazon Simple Notification Service (Amazon SNS)
- Amazon Simple Queue Service (Amazon SQS)
- AWS Step Functions

Compute

- Amazon EC2
- AWS Elastic Beanstalk
- AWS Lambda

Containers

- Amazon Elastic Container Registry (Amazon ECR)
- Amazon Elastic Container Service (Amazon ECS)
- Amazon Elastic Kubernetes Service (Amazon EKS)

Database

- Amazon Aurora
- Amazon DynamoDB
- Amazon ElastiCache
- Amazon RDS

Developer Tools

- AWS Amplify
- AWS Cloud9

- AWS CloudShell
- AWS CodeArtifact
- AWS CodeBuild
- AWS CodeCommit
- AWS CodeDeploy
- Amazon CodeGuru
- AWS CodePipeline
- AWS CodeStar
- AWS X-Ray

Management and Governance

- AWS AppConfig
- AWS Cloud Development Kit (AWS CDK)
- AWS CloudFormation
- AWS CloudTrail
- Amazon CloudWatch
- AWS Command Line Interface (AWS CLI)
- AWS Systems Manager

Networking and Content Delivery

- Amazon API Gateway
- Amazon CloudFront
- Elastic Load Balancing
- Amazon Route 53
- Amazon VPC

Security, Identity, and Compliance

- Amazon Cognito

- AWS Identity and Access Management (IAM)
- AWS Key Management Service (AWS KMS)
- AWS Secrets Manager
- AWS Security Token Service (AWS STS)
- AWS WAF

Storage

- Amazon Elastic Block Store (Amazon EBS)
- Amazon Elastic File System (Amazon EFS)
- Amazon S3

Out-of-Scope AWS Services

The following list contains AWS services and features that are out of scope for the AWS Certified Developer - Associate (DVA-C02) exam. This list is non-exhaustive and is subject to change.

Rubriques

- [Analytics](#)
- [Business Applications](#)
- [Compute](#)
- [Database](#)
- [End User Computing](#)
- [Internet of Things \(IoT\)](#)
- [Machine Learning](#)
- [Management and Governance](#)
- [Media Services](#)
- [Migration and Transfer](#)
- [Networking and Content Delivery](#)
- [Robotics](#)
- [Satellite](#)

- [Storage](#)

Analytics

- Amazon EMR
- AWS Glue
- Amazon Redshift

Business Applications

- Amazon Connect
- Amazon SES

Compute

- AWS Batch
- Amazon Lightsail
- AWS Outposts

Database

- Amazon DocumentDB
- Amazon Neptune
- Amazon Quantum Ledger Database (Amazon QLDB)

End User Computing

- Amazon AppStream 2.0
- Amazon WorkSpaces

Internet of Things (IoT)

- AWS IoT Core
- AWS IoT Greengrass

Machine Learning

- Amazon Comprehend
- Amazon Forecast
- Amazon Lex
- Amazon Polly
- Amazon Rekognition
- Amazon SageMaker
- Amazon Textract
- Amazon Transcribe
- Amazon Translate

Management and Governance

- AWS Config
- AWS Control Tower
- AWS License Manager
- AWS Organizations
- AWS Service Catalog
- AWS Trusted Advisor

Media Services

- Amazon Elastic Transcoder
- Amazon Kinesis Video Streams

Migration and Transfer

- AWS Database Migration Service (AWS DMS)
- AWS DataSync
- AWS Migration Hub
- AWS Snow Family

- AWS Transfer Family

Networking and Content Delivery

- AWS App Mesh
- AWS Cloud Map
- AWS Direct Connect
- AWS Global Accelerator
- AWS PrivateLink
- AWS Transit Gateway

Robotics

- AWS RoboMaker

Satellite

- AWS Ground Station

Storage

- AWS Backup
- Amazon FSx
- AWS Storage Gateway

AWS Certified SysOps Administrator - Associate (SOA-C02)

The AWS Certified SysOps Administrator - Associate (SOA-C02) exam is intended for system administrators in a cloud operations role. The exam validates a candidate's ability to deploy, manage, and operate workloads on AWS.

Rubriques

- [Introduction](#)
- [Target Candidate Description](#)

- [Exam Content](#)
- [Content Domains](#)
- [AWS Services for the Exam](#)
- [Domain 1: Monitoring, Logging, and Remediation \(20% of the exam content\)](#)
- [Domain 2: Reliability and Business Continuity \(16% of the exam content\)](#)
- [Domain 3: Deployment, Provisioning, and Automation \(18% of the exam content\)](#)
- [Domain 4: Security and Compliance \(16% of the exam content\)](#)
- [Domain 5: Networking and Content Delivery \(18% of the exam content\)](#)
- [Domain 6: Cost and Performance Optimization \(12% of the exam content\)](#)
- [In-Scope AWS Services](#)
- [Out-of-Scope AWS Services](#)

Introduction

The AWS Certified SysOps Administrator - Associate (SOA-C02) exam is intended for system administrators in a cloud operations role. The exam validates a candidate's ability to deploy, manage, and operate workloads on AWS.

The exam also validates a candidate's ability to complete the following tasks:

- Support and maintain AWS workloads according to the AWS Well-Architected Framework.
- Perform operations by using the AWS Management Console and the AWS CLI.
- Implement security controls to meet compliance requirements.
- Monitor, log, and troubleshoot systems.
- Apply networking concepts (for example, DNS, TCP/IP, firewalls).
- Implement architectural requirements (for example, high availability, performance, capacity).
- Perform business continuity and disaster recovery procedures.
- Identify, classify, and remediate incidents.

Target Candidate Description

The target candidate should have 1 year of experience with deployment, management, networking, and security on AWS.

The target candidate should have the following general IT knowledge and experience:

- 1–2 years of experience as a system administrator in an operations role
- Experience in monitoring, logging, and troubleshooting
- Knowledge of networking concepts (for example, DNS, TCP/IP, firewalls)
- Ability to implement architectural requirements (for example, high availability, performance, capacity)

The target candidate should have the following AWS knowledge and experience:

- Minimum of 1 year of hands-on experience with AWS technology
- Experience in deploying, managing, and operating workloads on AWS
- Understanding of the AWS Well-Architected Framework
- Hands-on experience with the AWS Management Console and the AWS CLI
- Understanding of AWS networking and security services
- Hands-on experience in implementing security controls and compliance requirements

Exam Content

The exam consists of two types of questions:

- Multiple choice: Has one correct response and three incorrect responses (distractors)
- Multiple response: Has two or more correct responses out of five or more response options

Select one or more responses that best complete the statement or answer the question. Distractors, or incorrect answers, are response options that a candidate with incomplete knowledge or skill might choose. Distractors are generally plausible responses that match the content area.

Unanswered questions are scored as incorrect; there is no penalty for guessing.

Content Domains

For detailed information about each domain, see the following sections:

- [the section called “Domain 1: Monitoring, Logging, and Remediation \(20% of the exam content\)”](#)
- [the section called “Domain 2: Reliability and Business Continuity \(16% of the exam content\)”](#)

- [the section called “Domain 3: Deployment, Provisioning, and Automation \(18% of the exam content\)”](#)
- [the section called “Domain 4: Security and Compliance \(16% of the exam content\)”](#)
- [the section called “Domain 5: Networking and Content Delivery \(18% of the exam content\)”](#)
- [the section called “Domain 6: Cost and Performance Optimization \(12% of the exam content\)”](#)

AWS Services for the Exam

The AWS Certified SysOps Administrator - Associate exam covers specific AWS services that are relevant to system administrators in cloud operations roles. Understanding which services are in scope and which are out of scope can help you focus your preparation efforts.

For detailed information about the AWS services covered in the exam, see the following sections:

- [In-Scope AWS Services](#)
- [Out-of-Scope AWS Services](#)

Domain 1: Monitoring, Logging, and Remediation (20% of the exam content)

This domain accounts for 20% of the exam content.

Rubriques

- [Task 1.1: Implement metrics, alarms, and filters by using monitoring and logging services](#)
- [Task 1.2: Remediate issues based on monitoring and availability metrics](#)

Task 1.1: Implement metrics, alarms, and filters by using monitoring and logging services

- Identify, collect, analyze, and export logs (for example, Amazon CloudWatch Logs, CloudWatch Logs Insights, CloudTrail logs).
- Collect metrics and logs by using the CloudWatch agent.
- Create CloudWatch alarms.
- Create metric filters.
- Create CloudWatch dashboards.

- Configure notifications (for example, Amazon Simple Notification Service [Amazon SNS], Service Quotas, CloudWatch alarms, Health events).

Task 1.2: Remediate issues based on monitoring and availability metrics

- Troubleshoot or take corrective actions based on notifications and alarms.
- Configure Amazon EventBridge rules to invoke actions.
- Use Systems Manager Automation runbooks to take action based on Config rules.

Domain 2: Reliability and Business Continuity (16% of the exam content)

This domain accounts for 16% of the exam content.

Rubriques

- [Task 2.1: Implement scalability and elasticity](#)
- [Task 2.2: Implement high availability and resilient environments](#)
- [Task 2.3: Implement backup and restore strategies](#)

Task 2.1: Implement scalability and elasticity

- Create and maintain Auto Scaling plans.
- Implement caching.
- Implement Amazon RDS replicas and Amazon Aurora Replicas.
- Implement loosely coupled architectures.
- Differentiate between horizontal scaling and vertical scaling.

Task 2.2: Implement high availability and resilient environments

- Configure Elastic Load Balancing (ELB) and Amazon Route 53 health checks.
- Differentiate between the use of a single Availability Zone and Multi-AZ deployments (for example, Amazon EC2 Auto Scaling groups, ELB, Amazon FSx, Amazon RDS).
- Implement fault-tolerant workloads (for example, Amazon Elastic File System [Amazon EFS], Elastic IP addresses).
- Implement Route 53 routing policies (for example, failover, weighted, latency based).

Task 2.3: Implement backup and restore strategies

- Automate snapshots and backups based on use cases (for example, RDS snapshots, Backup, RTO and RPO, Amazon Data Lifecycle Manager, retention policy).
- Restore databases (for example, point-in-time restore, promote read replica).
- Implement versioning and lifecycle rules.
- Configure Amazon S3 Cross-Region Replication (CRR).
- Perform disaster recovery procedures.

Domain 3: Deployment, Provisioning, and Automation (18% of the exam content)

This domain accounts for 18% of the exam content.

Rubriques

- [Task 3.1: Provision and maintain cloud resources](#)
- [Task 3.2: Automate manual or repeatable processes](#)

Task 3.1: Provision and maintain cloud resources

- Create and manage AMIs (for example, EC2 Image Builder).
- Create, manage, and troubleshoot CloudFormation.
- Provision resources across multiple Regions and accounts (for example, Resource Access Manager [RAM], CloudFormation StackSets, IAM cross-account roles).
- Select deployment scenarios and services (for example, blue/green, rolling, canary).
- Identify and remediate deployment issues (for example, service quotas, subnet sizing, CloudFormation errors, permissions).

Task 3.2: Automate manual or repeatable processes

- Use services (for example, Systems Manager, CloudFormation) to automate deployment processes.
- Implement automated patch management.
- Schedule automated tasks by using services (for example, EventBridge, Config).

Domain 4: Security and Compliance (16% of the exam content)

This domain accounts for 16% of the exam content.

Rubriques

- [Task 4.1: Implement and manage security and compliance policies](#)
- [Task 4.2: Implement data and infrastructure protection strategies](#)

Task 4.1: Implement and manage security and compliance policies

- Implement IAM features (for example, password policies, multi-factor authentication [MFA], roles, SAML, federated identity, resource policies, policy conditions).
- Troubleshoot and audit access issues by using services (for example, CloudTrail, IAM Access Analyzer, IAM policy simulator).
- Validate service control policies (SCPs) and permissions boundaries.
- Review Trusted Advisor security checks.
- Validate Region and service selections based on compliance requirements.
- Implement secure multi-account strategies (for example, Control Tower, Organizations).

Task 4.2: Implement data and infrastructure protection strategies

- Enforce a data classification scheme.
- Create, manage, and protect encryption keys.
- Implement encryption at rest (for example, Key Management Service [KMS]).
- Implement encryption in transit (for example, Certificate Manager [ACM], VPN).
- Securely store secrets by using services (for example, Secrets Manager, Systems Manager Parameter Store).
- Review reports or findings (for example, Security Hub, Amazon GuardDuty, Config, Amazon Inspector).

Domain 5: Networking and Content Delivery (18% of the exam content)

This domain accounts for 18% of the exam content.

Rubriques

- [Task 5.1: Implement networking features and connectivity](#)
- [Task 5.2: Configure domains, DNS services, and content delivery](#)
- [Task 5.3: Troubleshoot network connectivity issues](#)

Task 5.1: Implement networking features and connectivity

- Configure a VPC (for example, subnets, route tables, network ACLs, security groups, NAT gateway, internet gateway).
- Configure private connectivity (for example, Systems Manager Session Manager, VPC endpoints, VPC peering, VPN).
- Configure network protection services (for example, WAF, Shield).

Task 5.2: Configure domains, DNS services, and content delivery

- Configure Route 53 hosted zones and records.
- Implement Route 53 routing policies (for example, geolocation, geoproximity).
- Configure DNS (for example, Route 53 Resolver).
- Configure Amazon CloudFront and S3 origin access control (OAC).
- Configure S3 static website hosting.

Task 5.3: Troubleshoot network connectivity issues

- Interpret VPC configurations (for example, subnets, route tables, network ACLs, security groups).

Domain 6: Cost and Performance Optimization (12% of the exam content)

This domain accounts for 12% of the exam content.

Rubriques

- [Task 6.1: Implement cost optimization strategies](#)
- [Task 6.2: Implement performance optimization strategies](#)

Task 6.1: Implement cost optimization strategies

- Implement cost allocation tags.
- Identify and remediate underutilized or unused resources by using services and tools (for example, Trusted Advisor, Compute Optimizer, Cost Explorer).
- Configure Budgets and billing alarms.
- Assess resource usage patterns to qualify workloads for EC2 Spot Instances.
- Identify opportunities to use managed services (for example, Amazon RDS, Fargate, Amazon EFS).

Task 6.2: Implement performance optimization strategies

- Recommend compute resources based on performance metrics.
- Monitor Amazon Elastic Block Store (Amazon EBS) metrics and modify configuration to increase performance efficiency.
- Implement S3 performance features (for example, S3 Transfer Acceleration, multipart uploads).
- Monitor RDS metrics and modify the configuration to increase performance efficiency (for example, Performance Insights, RDS Proxy).
- Enable enhanced EC2 capabilities (for example, Elastic Network Adapter, instance store, placement groups).

In-Scope AWS Services

The following list contains AWS services and features that are in scope for the AWS Certified SysOps Administrator - Associate (SOA-C02) exam. This list is non-exhaustive and is subject to change. AWS offerings appear in categories that align with the offerings' primary functions.

Rubriques

- [Analytics](#)
- [Application Integration](#)
- [Cloud Financial Management](#)
- [Compute](#)
- [Database](#)
- [Developer Tools](#)

- [Management and Governance](#)
- [Migration and Transfer](#)
- [Networking and Content Delivery](#)
- [Security, Identity, and Compliance](#)
- [Storage](#)

Analytics

- Amazon OpenSearch Service

Application Integration

- Amazon EventBridge
- Amazon Simple Notification Service (Amazon SNS)
- Amazon Simple Queue Service (Amazon SQS)

Cloud Financial Management

- AWS Cost and Usage Report
- AWS Cost Explorer
- Savings Plans

Compute

- AWS Auto Scaling
- Amazon EC2
- Amazon EC2 Auto Scaling
- Amazon EC2 Image Builder
- AWS Lambda

Database

- Amazon Aurora

- Amazon DynamoDB
- Amazon ElastiCache
- Amazon RDS

Developer Tools

- AWS tools and SDKs

Management and Governance

- AWS CLI
- AWS CloudFormation
- AWS CloudTrail
- Amazon CloudWatch
- AWS Compute Optimizer
- AWS Config
- AWS Control Tower
- AWS Health Dashboard
- AWS License Manager
- AWS Management Console
- AWS Organizations
- AWS Service Catalog
- AWS Systems Manager
- AWS Trusted Advisor

Migration and Transfer

- AWS DataSync
- AWS Transfer Family

Networking and Content Delivery

- Amazon CloudFront

- Elastic Load Balancing (ELB)
- AWS Global Accelerator
- Amazon Route 53
- AWS Transit Gateway
- Amazon VPC
- AWS VPN

Security, Identity, and Compliance

- AWS Certificate Manager (ACM)
- Amazon Detective
- AWS Directory Service
- AWS Firewall Manager
- Amazon GuardDuty
- AWS Identity and Access Management (IAM)
- AWS Identity and Access Management Access Analyzer
- Amazon Inspector
- AWS Key Management Service (AWS KMS)
- AWS Secrets Manager
- AWS Security Hub
- AWS Shield
- AWS WAF

Storage

- AWS Backup
- Amazon Elastic Block Store (Amazon EBS)
- Amazon Elastic File System (Amazon EFS)
- Amazon FSx
- Amazon S3
- Amazon S3 Glacier
- AWS Storage Gateway

Out-of-Scope AWS Services

The following list contains AWS services and features that are out of scope for the AWS Certified SysOps Administrator - Associate (SOA-C02) exam. This list is non-exhaustive and is subject to change. AWS offerings that are entirely unrelated to the target job roles for the exam are excluded from this list.

Rubriques

- [Analytics](#)
- [Business Applications](#)
- [Compute](#)
- [Containers](#)
- [Database](#)
- [Developer Tools](#)
- [End User Computing](#)
- [Frontend Web and Mobile](#)
- [Game Tech](#)
- [Internet of Things \(IoT\)](#)
- [Machine Learning](#)
- [Management and Governance](#)
- [Media Services](#)
- [Migration and Transfer](#)
- [Security, Identity, and Compliance](#)
- [Storage](#)

Analytics

- Amazon EMR

Business Applications

- Amazon Chime
- Amazon Connect

- Amazon WorkDocs
- Amazon WorkMail

Compute

- Amazon Lightsail

Containers

- Amazon Elastic Container Registry (Amazon ECR)
- Amazon Elastic Container Service (Amazon ECS)

Database

- Amazon Redshift

Developer Tools

- AWS CodeBuild
- AWS CodeCommit
- AWS CodeDeploy
- AWS CodeStar
- AWS X-Ray

End User Computing

- Amazon AppStream 2.0
- Amazon WorkSpaces

Frontend Web and Mobile

- AWS Device Farm
- AWS Mobile SDKs
- Amazon Pinpoint

Game Tech

- Amazon GameLift

Internet of Things (IoT)

- AWS IoT Button
- AWS IoT Greengrass
- AWS IoT Platform

Machine Learning

- AWS Deep Learning AMIs (DLAMI)
- Amazon Lex
- Amazon Lumberyard
- Amazon Machine Learning (Amazon ML)
- Apache MXNet on AWS
- Amazon Polly
- Amazon Rekognition

Management and Governance

- AWS Managed Services (AMS)

Media Services

- Amazon Elastic Transcoder

Migration and Transfer

- AWS Schema Conversion Tool (AWS SCT)

Security, Identity, and Compliance

- Amazon Cloud Directory

Storage

- AWS Snowmobile

AWS Certified Solutions Architect - Associate (SAA-C03)

The AWS Certified Solutions Architect - Associate (SAA-C03) exam is intended for individuals who perform a solutions architect role. The exam validates a candidate's ability to design solutions based on the AWS Well-Architected Framework.

Rubriques

- [Introduction](#)
- [Target Candidate Description](#)
- [Exam Content](#)
- [Content Domains](#)
- [AWS Services for the Exam](#)
- [Domain 1: Design Secure Architectures \(30% of the exam content\)](#)
- [Domain 2: Design Resilient Architectures \(26% of the exam content\)](#)
- [Domain 3: Design High-Performing Architectures \(24% of the exam content\)](#)
- [Domain 4: Design Cost-Optimized Architectures \(20% of the exam content\)](#)
- [In-Scope AWS Services](#)
- [Out-of-Scope AWS Services](#)

Introduction

The AWS Certified Solutions Architect - Associate (SAA-C03) exam is intended for individuals who perform a solutions architect role. The exam validates a candidate's ability to design solutions based on the AWS Well-Architected Framework.

The exam also validates a candidate's ability to complete the following tasks:

- Design solutions that incorporate AWS services to meet current business requirements and future projected needs
- Design architectures that are secure, resilient, high-performing, and cost-optimized
- Review existing solutions and determine improvements

Target Candidate Description

The target candidate should have at least 1 year of hands-on experience designing cloud solutions that use AWS services.

Exam Content

There are two types of questions on the exam:

- Multiple choice: Has one correct response and three incorrect responses (distractors)
- Multiple response: Has two or more correct responses out of five or more response options

Select one or more responses that best complete the statement or answer the question. Distractors, or incorrect answers, are response options that a candidate with incomplete knowledge or skill might choose. Distractors are generally plausible responses that match the content area.

Unanswered questions are scored as incorrect; there is no penalty for guessing. The exam includes 50 questions that affect your score.

The exam includes 15 unscored questions that do not affect your score. AWS collects information about performance on these unscored questions to evaluate these questions for future use as scored questions. These unscored questions are not identified on the exam.

The AWS Certified Solutions Architect - Associate (SAA-C03) exam has a pass or fail designation. The exam is scored against a minimum standard established by AWS professionals who follow certification industry best practices and guidelines.

Your results for the exam are reported as a scaled score of 100–1,000. The minimum passing score is 720. Your score shows how you performed on the exam as a whole and whether you passed. Scaled scoring models help equate scores across multiple exam forms that might have slightly different difficulty levels.

Content Domains

For detailed information about each domain, see the following sections:

- [the section called “Domain 1: Design Secure Architectures \(30% of the exam content\)”](#)
- [the section called “Domain 2: Design Resilient Architectures \(26% of the exam content\)”](#)
- [the section called “Domain 3: Design High-Performing Architectures \(24% of the exam content\)”](#)
- [the section called “Domain 4: Design Cost-Optimized Architectures \(20% of the exam content\)”](#)

AWS Services for the Exam

The AWS Certified Solutions Architect - Associate exam covers specific AWS services that are relevant to designing solutions on AWS. Understanding which services are in scope and which are out of scope can help you focus your preparation efforts.

For detailed information about the AWS services covered in the exam, see the following sections:

- [In-Scope AWS Services](#)
- [Out-of-Scope AWS Services](#)

Domain 1: Design Secure Architectures (30% of the exam content)

This domain accounts for 30% of the exam content.

Rubriques

- [Task 1.1: Design secure access to resources](#)
- [Task 1.2: Design secure workloads and applications](#)
- [Task 1.3: Determine appropriate data security controls](#)

Task 1.1: Design secure access to resources

Knowledge of:

- Access controls and management across multiple accounts
- federated access and identity services (for example, Identity and Access Management [IAM], IAM Identity Center [Single Sign-On])

- global infrastructure (for example, Availability Zones, Regions)
- security best practices (for example, the principle of least privilege)
- The shared responsibility model

Skills in:

- Applying security best practices to IAM users and root users (for example, multi-factor authentication [MFA])
- Designing a flexible authorization model that includes IAM users, groups, roles, and policies
- Designing a role-based access control strategy (for example, Security Token Service [STS], role switching, cross-account access)
- Designing a security strategy for multiple accounts (for example, Control Tower, service control policies [SCPs])
- Determining the appropriate use of resource policies for services
- Determining when to federate a directory service with IAM roles

Task 1.2: Design secure workloads and applications

Knowledge of:

- Application configuration and credentials security
- service endpoints
- Control ports, protocols, and network traffic on
- Secure application access
- Security services with appropriate use cases (for example, Amazon Cognito, Amazon GuardDuty, Amazon Macie)
- Threat vectors external to (for example, DDoS, SQL injection)

Skills in:

- Designing VPC architectures with security components (for example, security groups, route tables, network ACLs, NAT gateways)
- Determining network segmentation strategies (for example, using public subnets and private subnets)

- Integrating services to secure applications (for example, Shield, WAF, IAM Identity Center, Secrets Manager)
- Securing external network connections to and from the Cloud (for example, VPN, Direct Connect)

Task 1.3: Determine appropriate data security controls

Knowledge of:

- Data access and governance
- Data recovery
- Data retention and classification
- Encryption and appropriate key management

Skills in:

- Aligning technologies to meet compliance requirements
- Encrypting data at rest (for example, Key Management Service [KMS])
- Encrypting data in transit (for example, Certificate Manager [ACM] using TLS)
- Implementing access policies for encryption keys
- Implementing data backups and replications
- Implementing policies for data access, lifecycle, and protection
- Rotating encryption keys and renewing certificates

Domain 2: Design Resilient Architectures (26% of the exam content)

This domain accounts for 26% of the exam content.

Rubriques

- [Task 2.1: Design scalable and loosely coupled architectures](#)
- [Task 2.2: Design highly available and/or fault-tolerant architectures](#)

Task 2.1: Design scalable and loosely coupled architectures

Knowledge of:

- API creation and management (for example, Amazon API Gateway, REST API)
- managed services with appropriate use cases (for example, Transfer Family, Amazon Simple Queue Service [Amazon SQS], Secrets Manager)
- Caching strategies
- Design principles for microservices (for example, stateless workloads compared with stateful workloads)
- Event-driven architectures
- Horizontal scaling and vertical scaling
- How to appropriately use edge accelerators (for example, content delivery network [CDN])
- How to migrate applications into containers
- Load balancing concepts (for example, Application Load Balancer)
- Multi-tier architectures
- Queuing and messaging concepts (for example, publish/subscribe)
- Serverless technologies and patterns (for example, Fargate, Lambda)
- Storage types with associated characteristics (for example, object, file, block)
- The orchestration of containers (for example, Amazon Elastic Container Service [Amazon ECS], Amazon Elastic Kubernetes Service [Amazon EKS])
- When to use read replicas
- Workflow orchestration (for example, Step Functions)

Skills in:

- Designing event-driven, microservice, and/or multi-tier architectures based on requirements
- Determining scaling strategies for components used in an architecture design
- Determining the services required to achieve loose coupling based on requirements
- Determining when to use containers
- Determining when to use serverless technologies and patterns
- Recommending appropriate compute, storage, networking, and database technologies based on requirements
- Using purpose-built services for workloads

Task 2.2: Design highly available and/or fault-tolerant architectures

Knowledge of:

- global infrastructure (for example, Availability Zones, Regions, Amazon Route 53)
- managed services with appropriate use cases (for example, Amazon Comprehend, Amazon Polly)
- Basic networking concepts (for example, route tables)
- Disaster recovery (DR) strategies (for example, backup and restore, pilot light, warm standby, active-active failover, recovery point objective [RPO], recovery time objective [RTO])
- Distributed design patterns
- Failover strategies
- Immutable infrastructure
- Load balancing concepts (for example, Application Load Balancer)
- Proxy concepts (for example, Amazon RDS Proxy)
- Service quotas and throttling (for example, how to configure the service quotas for a workload in a standby environment)
- Storage options and characteristics (for example, durability, replication)
- Workload visibility (for example, X-Ray)

Skills in:

- Determining automation strategies to ensure infrastructure integrity
- Determining the services required to provide a highly available and/or fault-tolerant architecture across Regions or Availability Zones
- Identifying metrics based on business requirements to deliver a highly available solution
- Implementing designs to mitigate single points of failure
- Implementing strategies to ensure the durability and availability of data (for example, backups)
- Selecting an appropriate DR strategy to meet business requirements
- Using services that improve the reliability of legacy applications and applications not built for the cloud (for example, when application changes are not possible)
- Using purpose-built services for workloads

Domain 3: Design High-Performing Architectures (24% of the exam content)

This domain accounts for 24% of the exam content.

Rubriques

- [Task 3.1: Determine high-performing and/or scalable storage solutions](#)
- [Task 3.2: Design high-performing and elastic compute solutions](#)
- [Task 3.3: Determine high-performing database solutions](#)
- [Task 3.4: Determine high-performing and/or scalable network architectures](#)
- [Task 3.5: Determine high-performing data ingestion and transformation solutions](#)

Task 3.1: Determine high-performing and/or scalable storage solutions

Knowledge of:

- Hybrid storage solutions to meet business requirements
- Storage services with appropriate use cases (for example, Amazon S3, Amazon Elastic File System [Amazon EFS], Amazon Elastic Block Store [Amazon EBS])
- Storage types with associated characteristics (for example, object, file, block)

Skills in:

- Determining storage services and configurations that meet performance demands
- Determining storage services that can scale to accommodate future needs

Task 3.2: Design high-performing and elastic compute solutions

Knowledge of:

- compute services with appropriate use cases (for example, Batch, Amazon EMR, Fargate)
- Distributed computing concepts supported by global infrastructure and edge services
- Queuing and messaging concepts (for example, publish/subscribe)
- Scalability capabilities with appropriate use cases (for example, Amazon EC2 Auto Scaling, Auto Scaling)

- Serverless technologies and patterns (for example, Lambda, Fargate)
- The orchestration of containers (for example, Amazon ECS, Amazon EKS)

Skills in:

- Decoupling workloads so that components can scale independently
- Identifying metrics and conditions to perform scaling actions
- Selecting the appropriate compute options and features (for example, EC2 instance types) to meet business requirements
- Selecting the appropriate resource type and size (for example, the amount of Lambda memory) to meet business requirements

Task 3.3: Determine high-performing database solutions

Knowledge of:

- global infrastructure (for example, Availability Zones, Regions)
- Caching strategies and services (for example, Amazon ElastiCache)
- Data access patterns (for example, read-intensive compared with write-intensive)
- Database capacity planning (for example, capacity units, instance types, Provisioned IOPS)
- Database connections and proxies
- Database engines with appropriate use cases (for example, heterogeneous migrations, homogeneous migrations)
- Database replication (for example, read replicas)
- Database types and services (for example, serverless, relational compared with non-relational, in-memory)

Skills in:

- Configuring read replicas to meet business requirements
- Designing database architectures
- Determining an appropriate database engine (for example, MySQL compared with PostgreSQL)
- Determining an appropriate database type (for example, Amazon Aurora, Amazon DynamoDB)

- Integrating caching to meet business requirements

Task 3.4: Determine high-performing and/or scalable network architectures

Knowledge of:

- Edge networking services with appropriate use cases (for example, Amazon CloudFront, Global Accelerator)
- How to design network architecture (for example, subnet tiers, routing, IP addressing)
- Load balancing concepts (for example, Application Load Balancer)
- Network connection options (for example, VPN, Direct Connect, PrivateLink)

Skills in:

- Creating a network topology for various architectures (for example, global, hybrid, multi-tier)
- Determining network configurations that can scale to accommodate future needs
- Determining the appropriate placement of resources to meet business requirements
- Selecting the appropriate load balancing strategy

Task 3.5: Determine high-performing data ingestion and transformation solutions

Knowledge of:

- Data analytics and visualization services with appropriate use cases (for example, Amazon Athena, Lake Formation, Amazon QuickSight)
- Data ingestion patterns (for example, frequency)
- Data transfer services with appropriate use cases (for example, DataSync, Storage Gateway)
- Data transformation services with appropriate use cases (for example, Glue)
- Secure access to ingestion access points
- Sizes and speeds needed to meet business requirements
- Streaming data services with appropriate use cases (for example, Amazon Kinesis)

Skills in:

- Building and securing data lakes

- Designing data streaming architectures
- Designing data transfer solutions
- Implementing visualization strategies
- Selecting appropriate compute options for data processing (for example, Amazon EMR)
- Selecting appropriate configurations for ingestion
- Transforming data between formats (for example, .csv to .parquet)

Domain 4: Design Cost-Optimized Architectures (20% of the exam content)

This domain accounts for 20% of the exam content.

Rubriques

- [Task 4.1: Design cost-optimized storage solutions](#)
- [Task 4.2: Design cost-optimized compute solutions](#)
- [Task 4.3: Design cost-optimized database solutions](#)
- [Task 4.4: Design cost-optimized network architectures](#)

Task 4.1: Design cost-optimized storage solutions

Knowledge of:

- Access options (for example, an S3 bucket with Requester Pays object storage)
- cost management service features (for example, cost allocation tags, multi-account billing)
- cost management tools with appropriate use cases (for example, Cost Explorer, Budgets, Cost and Usage Report)
- storage services with appropriate use cases (for example, Amazon FSx, Amazon EFS, Amazon S3, Amazon EBS)
- Backup strategies
- Block storage options (for example, hard disk drive [HDD] volume types, solid state drive [SSD] volume types)
- Data lifecycles
- Hybrid storage options (for example, DataSync, Transfer Family, Storage Gateway)

- Storage access patterns
- Storage tiering (for example, cold tiering for object storage)
- Storage types with associated characteristics (for example, object, file, block)

Skills in:

- Designing appropriate storage strategies (for example, batch uploads to Amazon S3 compared with individual uploads)
- Determining the correct storage size for a workload
- Determining the lowest cost method of transferring data for a workload to storage
- Determining when storage auto scaling is required
- Managing S3 object lifecycles
- Selecting the appropriate backup and/or archival solution
- Selecting the appropriate service for data migration to storage services
- Selecting the appropriate storage tier
- Selecting the correct data lifecycle for storage
- Selecting the most cost-effective storage service for a workload

Task 4.2: Design cost-optimized compute solutions

Knowledge of:

- cost management service features (for example, cost allocation tags, multi-account billing)
- cost management tools with appropriate use cases (for example, Cost Explorer, Budgets, Cost and Usage Report)
- global infrastructure (for example, Availability Zones, Regions)
- purchasing options (for example, Spot Instances, Reserved Instances, Savings Plans)
- Distributed compute strategies (for example, edge processing)
- Hybrid compute options (for example, Outposts, Snowball Edge)
- Instance types, families, and sizes (for example, memory optimized, compute optimized, virtualization)
- Optimization of compute utilization (for example, containers, serverless computing, microservices)

- Scaling strategies (for example, auto scaling, hibernation)

Skills in:

- Determining an appropriate load balancing strategy (for example, Application Load Balancer [Layer 7] compared with Network Load Balancer [Layer 4] compared with Gateway Load Balancer)
- Determining appropriate scaling methods and strategies for elastic workloads (for example, horizontal compared with vertical, EC2 hibernation)
- Determining cost-effective compute services with appropriate use cases (for example, Lambda, Amazon EC2, Fargate)
- Determining the required availability for different classes of workloads (for example, production workloads, non-production workloads)
- Selecting the appropriate instance family for a workload
- Selecting the appropriate instance size for a workload

Task 4.3: Design cost-optimized database solutions

Knowledge of:

- cost management service features (for example, cost allocation tags, multi-account billing)
- cost management tools with appropriate use cases (for example, Cost Explorer, Budgets, Cost and Usage Report)
- Caching strategies
- Data retention policies
- Database capacity planning (for example, capacity units)
- Database connections and proxies
- Database engines with appropriate use cases (for example, heterogeneous migrations, homogeneous migrations)
- Database replication (for example, read replicas)
- Database types and services (for example, relational compared with non-relational, Aurora, DynamoDB)

Skills in:

- Designing appropriate backup and retention policies (for example, snapshot frequency)
- Determining an appropriate database engine (for example, MySQL compared with PostgreSQL)
- Determining cost-effective database services with appropriate use cases (for example, DynamoDB compared with Amazon RDS, serverless)
- Determining cost-effective database types (for example, time series format, columnar format)
- Migrating database schemas and data to different locations and/or different database engines

Task 4.4: Design cost-optimized network architectures

Knowledge of:

- cost management service features (for example, cost allocation tags, multi-account billing)
- cost management tools with appropriate use cases (for example, Cost Explorer, Budgets, Cost and Usage Report)
- Load balancing concepts (for example, Application Load Balancer)
- NAT gateways (for example, NAT instance costs compared with NAT gateway costs)
- Network connectivity (for example, private lines, dedicated lines, VPNs)
- Network routing, topology, and peering (for example, Transit Gateway, VPC peering)
- Network services with appropriate use cases (for example, DNS)

Skills in:

- Configuring appropriate NAT gateway types for a network (for example, a single shared NAT gateway compared with NAT gateways for each Availability Zone)
- Configuring appropriate network connections (for example, Direct Connect compared with VPN compared with internet)
- Configuring appropriate network routes to minimize network transfer costs (for example, Region to Region, Availability Zone to Availability Zone, private to public, Global Accelerator, VPC endpoints)
- Determining strategic needs for content delivery networks (CDNs) and edge caching
- Reviewing existing workloads for network optimizations
- Selecting an appropriate throttling strategy
- Selecting the appropriate bandwidth allocation for a network device (for example, a single VPN compared with multiple VPNs, Direct Connect speed)

In-Scope AWS Services

The following list contains AWS services and features that are in scope for the AWS Certified Solutions Architect - Associate (SAA-C03) exam. This list is non-exhaustive and is subject to change. AWS offerings appear in categories that align with the offerings' primary functions.

Rubriques

- [Analytics](#)
- [Application Integration](#)
- [Business Applications](#)
- [Compute](#)
- [Containers](#)
- [Database](#)
- [Developer Tools](#)
- [Frontend Web and Mobile](#)
- [Management and Governance](#)
- [Migration and Transfer](#)
- [Networking and Content Delivery](#)
- [Security, Identity, and Compliance](#)
- [Storage](#)

Analytics

- Amazon Athena
- AWS Data Exchange
- AWS Data Pipeline
- Amazon EMR
- AWS Glue
- Amazon Kinesis
- AWS Lake Formation
- Amazon Managed Streaming for Apache Kafka (Amazon MSK)
- Amazon OpenSearch Service

- Amazon QuickSight
- Amazon Redshift

Application Integration

- Amazon AppFlow
- AWS AppSync
- Amazon EventBridge
- Amazon MQ
- Amazon Simple Notification Service (Amazon SNS)
- Amazon Simple Queue Service (Amazon SQS)
- AWS Step Functions

Business Applications

- Amazon Simple Email Service (Amazon SES)

Compute

- AWS Batch
- Amazon EC2
- Amazon EC2 Auto Scaling
- AWS Elastic Beanstalk
- AWS Lambda
- Amazon Lightsail
- AWS Outposts
- AWS Serverless Application Repository
- AWS Wavelength

Containers

- Amazon Elastic Container Registry (Amazon ECR)

- Amazon Elastic Container Service (Amazon ECS)
- Amazon Elastic Kubernetes Service (Amazon EKS)
- AWS Fargate

Database

- Amazon Aurora
- Amazon DocumentDB (with MongoDB compatibility)
- Amazon DynamoDB
- Amazon ElastiCache
- Amazon Keyspaces (for Apache Cassandra)
- Amazon Neptune
- Amazon RDS
- Amazon Redshift

Developer Tools

- AWS X-Ray

Frontend Web and Mobile

- AWS Amplify
- AWS AppSync
- AWS Device Farm

Management and Governance

- AWS Auto Scaling
- AWS Backup
- AWS CloudFormation
- AWS CloudTrail
- Amazon CloudWatch
- AWS Compute Optimizer

- AWS Config
- AWS Control Tower
- Amazon Data Lifecycle Manager
- AWS License Manager
- AWS Organizations
- AWS Resource Access Manager
- AWS Service Catalog
- AWS Systems Manager
- AWS Trusted Advisor
- AWS Well-Architected Tool

Migration and Transfer

- AWS Application Discovery Service
- AWS Application Migration Service
- AWS Database Migration Service (AWS DMS)
- AWS DataSync
- AWS Migration Hub
- AWS Schema Conversion Tool (AWS SCT)
- AWS Snow Family
- AWS Transfer Family

Networking and Content Delivery

- Amazon API Gateway
- AWS Client VPN
- Amazon CloudFront
- AWS Direct Connect
- Elastic Load Balancing
- AWS Global Accelerator
- AWS PrivateLink

- Amazon Route 53
- AWS Site-to-Site VPN
- AWS Transit Gateway
- Amazon VPC
- AWS VPN

Security, Identity, and Compliance

- AWS Certificate Manager (ACM)
- AWS CloudHSM
- Amazon Cognito
- Amazon Detective
- AWS Directory Service
- AWS Firewall Manager
- Amazon GuardDuty
- AWS Identity and Access Management (IAM)
- Amazon Inspector
- AWS Key Management Service (AWS KMS)
- Amazon Macie
- AWS Network Firewall
- AWS Resource Access Manager (AWS RAM)
- AWS Secrets Manager
- AWS Security Hub
- AWS Shield
- AWS Single Sign-On
- AWS WAF

Storage

- AWS Backup
- Amazon Elastic Block Store (Amazon EBS)

- Amazon Elastic File System (Amazon EFS)
- Amazon FSx (for all types)
- Amazon S3
- Amazon S3 Glacier
- AWS Storage Gateway

Out-of-Scope AWS Services

The following list contains AWS services and features that are out of scope for the AWS Certified Solutions Architect - Associate (SAA-C03) exam. This list is non-exhaustive and is subject to change.

Rubriques

- [Analytics](#)
- [Application Integration](#)
- [Business Applications](#)
- [Compute](#)
- [Developer Tools](#)
- [Game Tech](#)
- [Internet of Things \(IoT\)](#)
- [Machine Learning](#)
- [Media Services](#)
- [Robotics](#)
- [Satellite](#)

Analytics

- Amazon CloudSearch
- Amazon Managed Service for Apache Flink

Application Integration

- Amazon Simple Workflow Service

Business Applications

- Alexa for Business
- Amazon Chime
- Amazon Connect
- Amazon WorkDocs
- Amazon WorkMail

Compute

- Amazon AppStream 2.0
- AWS Mainframe Modernization
- Amazon WorkSpaces

Developer Tools

- AWS Cloud9
- AWS CodeBuild
- AWS CodeCommit
- AWS CodeDeploy
- Amazon CodeGuru
- AWS CodePipeline
- AWS CodeStar

Game Tech

- Amazon GameLift
- Amazon Lumberyard

Internet of Things (IoT)

- AWS IoT 1-Click
- AWS IoT Analytics

- AWS IoT Button
- AWS IoT Core
- AWS IoT Device Defender
- AWS IoT Device Management
- AWS IoT Events
- AWS IoT Greengrass
- AWS IoT SiteWise
- AWS IoT Things Graph

Machine Learning

- Amazon Augmented AI
- Amazon CodeGuru
- Amazon Comprehend
- AWS DeepComposer
- AWS DeepLens
- AWS DeepRacer
- Amazon DevOps Guru
- Amazon Forecast
- Amazon Fraud Detector
- Amazon Kendra
- Amazon Lex
- Amazon Lookout for Equipment
- Amazon Lookout for Metrics
- Amazon Lookout for Vision
- Amazon Monitron
- Amazon Personalize
- Amazon Polly
- Amazon Rekognition
- Amazon SageMaker

- Amazon Textract
- Amazon Transcribe
- Amazon Translate

Media Services

- Amazon Elastic Transcoder
- Amazon Interactive Video Service
- Amazon Kinesis Video Streams
- AWS Elemental MediaConnect
- AWS Elemental MediaConvert
- AWS Elemental MediaLive
- AWS Elemental MediaPackage
- AWS Elemental MediaStore
- AWS Elemental MediaTailor

Robotics

- AWS RoboMaker

Satellite

- AWS Ground Station

AWS Certified Machine Learning Engineer - Associate (MLA-C01)

The AWS Certified Machine Learning Engineer - Associate (MLA-C01) exam validates a candidate's ability to build, operationalize, deploy, and maintain machine learning (ML) solutions and pipelines by using the AWS Cloud.

Rubriques

- [Introduction](#)
- [Target Candidate Description](#)

- [Exam Content](#)
- [Content Domains](#)
- [AWS Services for the Exam](#)
- [Domain 1: Data Preparation for Machine Learning \(ML\) \(28% of the exam content\)](#)
- [Domain 2: ML Model Development \(26% of the exam content\)](#)
- [Domain 3: Deployment and Orchestration of ML Workflows \(22% of the exam content\)](#)
- [Domain 4: ML Solution Monitoring, Maintenance, and Security \(24% of the exam content\)](#)
- [In-Scope AWS Services](#)

Introduction

The AWS Certified Machine Learning Engineer - Associate (MLA-C01) exam validates a candidate's ability to build, operationalize, deploy, and maintain machine learning (ML) solutions and pipelines by using the AWS Cloud.

The exam also validates a candidate's ability to complete the following tasks:

- Ingest, transform, validate, and prepare data for ML modeling.
- Select general modeling approaches, train models, tune hyperparameters, analyze model performance, and manage model versions.
- Choose deployment infrastructure and endpoints, provision compute resources, and configure auto scaling based on requirements.
- Set up continuous integration and continuous delivery (CI/CD) pipelines to automate orchestration of ML workflows.
- Monitor models, data, and infrastructure to detect issues.
- Secure ML systems and resources through access controls, compliance features, and best practices.

Target Candidate Description

The target candidate should have at least 1 year of experience using Amazon SageMaker and other AWS services for ML engineering. The target candidate also should have at least 1 year of experience in a related role such as a backend software developer, DevOps developer, data engineer, or data scientist.

The target candidate should have the following general IT knowledge:

- Basic understanding of common ML algorithms and their use cases
- Data engineering fundamentals, including knowledge of common data formats, ingestion, and transformation to work with ML data pipelines
- Knowledge of querying and transforming data
- Knowledge of software engineering best practices for modular, reusable code development, deployment, and debugging
- Familiarity with provisioning and monitoring cloud and on-premises ML resources
- Experience with CI/CD pipelines and infrastructure as code (IaC)
- Experience with code repositories for version control and CI/CD pipelines

The target candidate should have the following AWS knowledge:

- Knowledge of SageMaker capabilities and algorithms for model building and deployment
- Knowledge of AWS services for data storage, processing, and visualization
- Understanding of AWS security services and best practices
- Familiarity with AWS deployment and monitoring tools
- Experience with AWS infrastructure management and scaling

The following list contains job tasks that the target candidate is not expected to be able to perform. This list is non-exhaustive. These tasks are out of scope for the exam:

- Develop novel ML algorithms
- Perform advanced statistical analysis
- Design complex neural network architectures
- Implement custom ML frameworks
- Develop specialized ML hardware solutions

Exam Content

The exam consists of multiple-choice and multiple-response questions. Multiple-choice questions have one correct response and three incorrect responses (distractors). Multiple-response questions have two or more correct responses out of five or more options.

Select one or more responses that best complete the statement or answer the question. Distractors, or incorrect answers, are response options that a candidate with incomplete knowledge or skill might choose. Distractors are generally plausible responses that match the content area.

Unanswered questions are scored as incorrect; there is no penalty for guessing.

The exam includes unscored questions that do not affect your score. AWS collects information about performance on these unscored questions to evaluate these questions for future use as scored questions. These unscored questions are not identified on the exam.

The AWS Certified Machine Learning Engineer - Associate (MLA-C01) exam has a pass or fail designation. The exam is scored against a minimum standard established by AWS professionals who follow certification industry best practices and guidelines.

Your results for the exam are reported as a scaled score of 100–1,000. The minimum passing score is 750. Your score shows how you performed on the exam as a whole and whether you passed. Scaled scoring models help equate scores across multiple exam forms that might have slightly different difficulty levels.

Content Domains

For detailed information about each domain, see the following sections:

- [the section called “Domain 1: Data Preparation for Machine Learning \(ML\) \(28% of the exam content\)”](#)
- [the section called “Domain 2: ML Model Development \(26% of the exam content\)”](#)
- [the section called “Domain 3: Deployment and Orchestration of ML Workflows \(22% of the exam content\)”](#)
- [the section called “Domain 4: ML Solution Monitoring, Maintenance, and Security \(24% of the exam content\)”](#)

AWS Services for the Exam

The AWS Certified Machine Learning Engineer - Associate exam covers specific AWS services that are relevant to machine learning engineers. Understanding which services are in scope can help you focus your preparation efforts.

For detailed information about the AWS services covered in the exam, see the following section:

- [In-Scope AWS Services](#)

Domain 1: Data Preparation for Machine Learning (ML) (28% of the exam content)

This domain accounts for 28% of the exam content.

Rubriques

- [Task 1.1: Ingest and store data](#)
- [Task 1.2: Transform data and perform feature engineering](#)
- [Task 1.3: Ensure data integrity and prepare data for modeling](#)

Task 1.1: Ingest and store data

Knowledge of:

- Data formats and ingestion mechanisms (for example, validated and non-validated formats, Apache Parquet, JSON, CSV, Apache ORC, Apache Avro, RecordIO)
- How to use the core data sources (for example, Amazon S3, Amazon Elastic File System [Amazon EFS], Amazon FSx for NetApp ONTAP)
- How to use streaming data sources to ingest data (for example, Amazon Kinesis, Apache Flink, Apache Kafka)
- storage options, including use cases and tradeoffs

Skills in:

- Extracting data from storage (for example, Amazon S3, Amazon Elastic Block Store [Amazon EBS], Amazon EFS, Amazon RDS, Amazon DynamoDB) by using relevant service options (for example, Amazon S3 Transfer Acceleration, Amazon EBS Provisioned IOPS)
- Choosing appropriate data formats (for example, Parquet, JSON, CSV, ORC) based on data access patterns
- Ingesting data into Amazon SageMaker Data Wrangler and SageMaker Feature Store
- Merging data from multiple sources (for example, by using programming techniques, Glue, Apache Spark)

- Troubleshooting and debugging data ingestion and storage issues that involve capacity and scalability
- Making initial storage decisions based on cost, performance, and data structure

Task 1.2: Transform data and perform feature engineering

Knowledge of:

- Data cleaning and transformation techniques (for example, detecting and treating outliers, imputing missing data, combining, deduplication)
- Feature engineering techniques (for example, data scaling and standardization, feature splitting, binning, log transformation, normalization)
- Encoding techniques (for example, one-hot encoding, binary encoding, label encoding, tokenization)
- Tools to explore, visualize, or transform data and features (for example, SageMaker Data Wrangler, Glue, Glue DataBrew)
- Services that transform streaming data (for example, Lambda, Spark)
- Data annotation and labeling services that create high-quality labeled datasets

Skills in:

- Transforming data by using tools (for example, Glue, Glue DataBrew, Spark running on Amazon EMR, SageMaker Data Wrangler)
- Creating and managing features by using tools (for example, SageMaker Feature Store)
- Validating and labeling data by using services (for example, SageMaker Ground Truth, Amazon Mechanical Turk)

Task 1.3: Ensure data integrity and prepare data for modeling

Knowledge of:

- Pre-training bias metrics for numeric, text, and image data (for example, class imbalance [CI], difference in proportions of labels [DPL])
- Strategies to address CI in numeric, text, and image datasets (for example, synthetic data generation, resampling)

- Techniques to encrypt data
- Data classification, anonymization, and masking
- Implications of compliance requirements (for example, personally identifiable information [PII], protected health information [PHI], data residency)

Skills in:

- Validating data quality (for example, by using Glue DataBrew and Glue Data Quality)
- Identifying and mitigating sources of bias in data (for example, selection bias, measurement bias) by using tools (for example, SageMaker Clarify)
- Preparing data to reduce prediction bias (for example, by using dataset splitting, shuffling, and augmentation)
- Configuring data to load into the model training resource (for example, Amazon EFS, Amazon FSx)

Domain 2: ML Model Development (26% of the exam content)

This domain accounts for 26% of the exam content.

Rubriques

- [Task 2.1: Choose a modeling approach](#)
- [Task 2.2: Train and refine models](#)
- [Task 2.3: Analyze model performance](#)

Task 2.1: Choose a modeling approach

Knowledge of:

- Capabilities and appropriate uses of ML algorithms to solve business problems
- How to use artificial intelligence (AI) services (for example, Amazon Translate, Amazon Transcribe, Amazon Rekognition, Amazon Bedrock) to solve specific business problems
- How to consider interpretability during model selection or algorithm selection
- SageMaker built-in algorithms and when to apply them

Skills in:

- Assessing available data and problem complexity to determine the feasibility of an ML solution
- Comparing and selecting appropriate ML models or algorithms to solve specific problems
- Choosing built-in algorithms, foundation models, and solution templates (for example, in SageMaker JumpStart and Amazon Bedrock)
- Selecting models or algorithms based on costs
- Selecting AI services to solve common business needs

Task 2.2: Train and refine models

Knowledge of:

- Elements in the training process (for example, epoch, steps, batch size)
- Methods to reduce model training time (for example, early stopping, distributed training)
- Factors that influence model size
- Methods to improve model performance
- Benefits of regularization techniques (for example, dropout, weight decay, L1 and L2)
- Hyperparameter tuning techniques (for example, random search, Bayesian optimization)
- Model hyperparameters and their effects on model performance (for example, number of trees in a tree-based model, number of layers in a neural network)
- Methods to integrate models that were built outside SageMaker into SageMaker

Skills in:

- Using SageMaker built-in algorithms and common ML libraries to develop ML models
- Using SageMaker script mode with SageMaker supported frameworks to train models (for example, TensorFlow, PyTorch)
- Using custom datasets to fine-tune pre-trained models (for example, Amazon Bedrock, SageMaker JumpStart)
- Performing hyperparameter tuning (for example, by using SageMaker automatic model tuning [AMT])
- Integrating automated hyperparameter optimization capabilities
- Preventing model overfitting, underfitting, and catastrophic forgetting (for example, by using regularization techniques, feature selection)

- Combining multiple training models to improve performance (for example, ensembling, stacking, boosting)
- Reducing model size (for example, by altering data types, pruning, updating feature selection, compression)
- Managing model versions for repeatability and audits (for example, by using the SageMaker Model Registry)

Task 2.3: Analyze model performance

Knowledge of:

- Model evaluation techniques and metrics (for example, confusion matrix, heat maps, F1 score, accuracy, precision, recall, Root Mean Square Error [RMSE], receiver operating characteristic [ROC], Area Under the ROC Curve [AUC])
- Methods to create performance baselines
- Methods to identify model overfitting and underfitting
- Metrics available in SageMaker Clarify to gain insights into ML training data and models
- Convergence issues

Skills in:

- Selecting and interpreting evaluation metrics and detecting model bias
- Assessing tradeoffs between model performance, training time, and cost
- Performing reproducible experiments by using services
- Comparing the performance of a shadow variant to the performance of a production variant
- Using SageMaker Clarify to interpret model outputs
- Using SageMaker Model Debugger to debug model convergence

Domain 3: Deployment and Orchestration of ML Workflows (22% of the exam content)

This domain accounts for 22% of the exam content.

Rubriques

- [Task 3.1: Select deployment infrastructure based on existing architecture and requirements](#)
- [Task 3.2: Create and script infrastructure based on existing architecture and requirements](#)
- [Task 3.3: Use automated orchestration tools to set up continuous integration and continuous delivery \(CI/CD\) pipelines](#)

Task 3.1: Select deployment infrastructure based on existing architecture and requirements

Knowledge of:

- Deployment best practices (for example, versioning, rollback strategies)
- deployment services (for example, SageMaker)
- Methods to serve ML models in real time and in batches
- How to provision compute resources in production environments and test environments (for example, CPU, GPU)
- Model and endpoint requirements for deployment endpoints (for example, serverless endpoints, real-time endpoints, asynchronous endpoints, batch inference)
- How to choose appropriate containers (for example, provided or customized)
- Methods to optimize models on edge devices (for example, SageMaker Neo)

Skills in:

- Evaluating performance, cost, and latency tradeoffs
- Choosing the appropriate compute environment for training and inference based on requirements (for example, GPU or CPU specifications, processor family, networking bandwidth)
- Selecting the correct deployment orchestrator (for example, Apache Airflow, SageMaker Pipelines)
- Selecting multi-model or multi-container deployments
- Selecting the correct deployment target (for example, SageMaker endpoints, Kubernetes, Amazon Elastic Container Service [Amazon ECS], Amazon Elastic Kubernetes Service [Amazon EKS], Lambda)
- Choosing model deployment strategies (for example, real time, batch)

Task 3.2: Create and script infrastructure based on existing architecture and requirements

Knowledge of:

- Difference between on-demand and provisioned resources
- How to compare scaling policies
- Tradeoffs and use cases of infrastructure as code (IaC) options (for example, CloudFormation, Cloud Development Kit [CDK])
- Containerization concepts and container services
- How to use SageMaker endpoint auto scaling policies to meet scalability requirements (for example, based on demand, time)

Skills in:

- Applying best practices to enable maintainable, scalable, and cost-effective ML solutions (for example, automatic scaling on SageMaker endpoints, dynamically adding Spot Instances, by using Amazon EC2 instances, by using Lambda behind the endpoints)
- Automating the provisioning of compute resources, including communication between stacks (for example, by using CloudFormation, CDK)
- Building and maintaining containers (for example, Amazon Elastic Container Registry [Amazon ECR], Amazon EKS, Amazon ECS, by using bring your own container [BYOC] with SageMaker)
- Configuring SageMaker endpoints within the VPC network
- Deploying and hosting models by using the SageMaker SDK
- Choosing specific metrics for auto scaling (for example, model latency, CPU utilization, invocations per instance)

Task 3.3: Use automated orchestration tools to set up continuous integration and continuous delivery (CI/CD) pipelines

Knowledge of:

- Capabilities and quotas for CodePipeline, CodeBuild, and CodeDeploy
- Automation and integration of data ingestion with orchestration services

- Version control systems and basic usage (for example, Git)
- CI/CD principles and how they fit into ML workflows
- Deployment strategies and rollback actions (for example, blue/green, canary, linear)
- How code repositories and pipelines work together

Skills in:

- Configuring and troubleshooting CodeBuild, CodeDeploy, and CodePipeline, including stages
- Applying continuous deployment flow structures to invoke pipelines (for example, Gitflow, GitHub Flow)
- Using services to automate orchestration (for example, to deploy ML models, automate model building)
- Configuring training and inference jobs (for example, by using Amazon EventBridge rules, SageMaker Pipelines, CodePipeline)
- Creating automated tests in CI/CD pipelines (for example, integration tests, unit tests, end-to-end tests)
- Building and integrating mechanisms to retrain models

Domain 4: ML Solution Monitoring, Maintenance, and Security (24% of the exam content)

This domain accounts for 24% of the exam content.

Rubriques

- [Task 4.1: Monitor model inference](#)
- [Task 4.2: Monitor and optimize infrastructure and costs](#)
- [Task 4.3: Secure resources](#)

Task 4.1: Monitor model inference

Knowledge of:

- Drift in ML models

- Techniques to monitor data quality and model performance
- Design principles for ML lenses relevant to monitoring

Skills in:

- Monitoring models in production (for example, by using SageMaker Model Monitor)
- Monitoring workflows to detect anomalies or errors in data processing or model inference
- Detecting changes in the distribution of data that can affect model performance (for example, by using SageMaker Clarify)
- Monitoring model performance in production by using A/B testing

Task 4.2: Monitor and optimize infrastructure and costs

Knowledge of:

- Key performance metrics for ML infrastructure (for example, utilization, throughput, availability, scalability, fault tolerance)
- Monitoring and observability tools to troubleshoot latency and performance issues (for example, X-Ray, Amazon CloudWatch Lambda Insights, Amazon CloudWatch Logs Insights)
- How to use CloudTrail to log, monitor, and invoke re-training activities
- Differences between instance types and how they affect performance (for example, memory optimized, compute optimized, general purpose, inference optimized)
- Capabilities of cost analysis tools (for example, Cost Explorer, Billing and Cost Management, Trusted Advisor)
- Cost tracking and allocation techniques (for example, resource tagging)

Skills in:

- Configuring and using tools to troubleshoot and analyze resources (for example, CloudWatch Logs, CloudWatch alarms)
- Creating CloudTrail trails
- Setting up dashboards to monitor performance metrics (for example, by using Amazon QuickSight, CloudWatch dashboards)
- Monitoring infrastructure (for example, by using EventBridge events)

- Rightsizing instance families and sizes (for example, by using SageMaker Inference Recommender and Compute Optimizer)
- Monitoring and resolving latency and scaling issues
- Preparing infrastructure for cost monitoring (for example, by applying a tagging strategy)
- Troubleshooting capacity concerns that involve cost and performance (for example, provisioned concurrency, service quotas, auto scaling)
- Optimizing costs and setting cost quotas by using appropriate cost management tools (for example, Cost Explorer, Trusted Advisor, Budgets)
- Optimizing infrastructure costs by selecting purchasing options (for example, Spot Instances, On-Demand Instances, Reserved Instances, SageMaker Savings Plans)

Task 4.3: Secure resources

Knowledge of:

- IAM roles, policies, and groups that control access to services (for example, Identity and Access Management [IAM], bucket policies, SageMaker Role Manager)
- SageMaker security and compliance features
- Controls for network access to ML resources
- Security best practices for CI/CD pipelines

Skills in:

- Configuring least privilege access to ML artifacts
- Configuring IAM policies and roles for users and applications that interact with ML systems
- Monitoring, auditing, and logging ML systems to ensure continued security and compliance
- Troubleshooting and debugging security issues
- Building VPCs, subnets, and security groups to securely isolate ML systems

In-Scope AWS Services

The following list contains AWS services and features that are in scope for the AWS Certified Machine Learning Engineer - Associate (MLA-C01) exam. This list is non-exhaustive and is subject to change. AWS offerings appear in categories that align with the offerings' primary functions.

Rubriques

- [Analytics](#)
- [Application Integration](#)
- [Cloud Financial Management](#)
- [Compute](#)
- [Containers](#)
- [Database](#)
- [Developer Tools](#)
- [Machine Learning](#)
- [Management and Governance](#)
- [Media](#)
- [Migration and Transfer](#)
- [Networking and Content Delivery](#)
- [Security, Identity, and Compliance](#)
- [Storage](#)

Analytics

- Amazon Athena
- Amazon Data Firehose
- Amazon EMR
- AWS Glue
- AWS Glue DataBrew
- AWS Glue Data Quality
- Amazon Kinesis
- AWS Lake Formation
- Amazon Managed Service for Apache Flink
- Amazon OpenSearch Service
- Amazon QuickSight
- Amazon Redshift

Application Integration

- Amazon EventBridge
- Amazon Managed Workflows for Apache Airflow (Amazon MWAA)
- Amazon Simple Notification Service (Amazon SNS)
- Amazon Simple Queue Service (Amazon SQS)
- AWS Step Functions

Cloud Financial Management

- AWS Billing and Cost Management
- AWS Budgets
- AWS Cost Explorer

Compute

- AWS Batch
- Amazon EC2
- AWS Lambda
- AWS Serverless Application Repository

Containers

- Amazon Elastic Container Registry (Amazon ECR)
- Amazon Elastic Container Service (Amazon ECS)
- Amazon Elastic Kubernetes Service (Amazon EKS)

Database

- Amazon DocumentDB (with MongoDB compatibility)
- Amazon DynamoDB
- Amazon ElastiCache
- Amazon Neptune

- Amazon RDS

Developer Tools

- AWS Cloud Development Kit (AWS CDK)
- AWS CodeArtifact
- AWS CodeBuild
- AWS CodeDeploy
- AWS CodePipeline
- AWS X-Ray

Machine Learning

- Amazon Augmented AI (Amazon A2I)
- Amazon Bedrock
- Amazon CodeGuru
- Amazon Comprehend
- Amazon Comprehend Medical
- Amazon DevOps Guru
- Amazon Fraud Detector
- AWS HealthLake
- Amazon Kendra
- Amazon Lex
- Amazon Lookout for Equipment
- Amazon Lookout for Metrics
- Amazon Lookout for Vision
- Amazon Mechanical Turk
- Amazon Personalize
- Amazon Polly
- Amazon Q
- Amazon Rekognition

- Amazon SageMaker
- Amazon Textract
- Amazon Transcribe
- Amazon Translate

Management and Governance

- AWS Auto Scaling
- AWS Chatbot
- AWS CloudFormation
- AWS CloudTrail
- Amazon CloudWatch
- Amazon CloudWatch Logs
- AWS Compute Optimizer
- AWS Config
- AWS Organizations
- AWS Service Catalog
- AWS Systems Manager
- AWS Trusted Advisor

Media

- Amazon Kinesis Video Streams

Migration and Transfer

- AWS DataSync

Networking and Content Delivery

- Amazon API Gateway
- Amazon CloudFront

- AWS Direct Connect
- Amazon VPC

Security, Identity, and Compliance

- AWS Identity and Access Management (IAM)
- AWS Key Management Service (AWS KMS)
- Amazon Macie
- AWS Secrets Manager

Storage

- Amazon Elastic Block Store (Amazon EBS)
- Amazon Elastic File System (Amazon EFS)
- Amazon FSx
- Amazon S3
- Amazon S3 Glacier
- AWS Storage Gateway

AWS Certified Data Engineer - Associate (DEA-C01)

The AWS Certified Data Engineer - Associate (DEA-C01) exam validates a candidate's ability to implement data pipelines and to monitor, troubleshoot, and optimize cost and performance issues in accordance with best practices.

Rubriques

- [Introduction](#)
- [Target Candidate Description](#)
- [Exam Content](#)
- [Content Domains](#)
- [AWS Services for the Exam](#)
- [Domain 1: Data Ingestion and Transformation \(36% of the exam content\)](#)

- [Domain 2: Data Store Management \(26% of the exam content\)](#)
- [Domain 3: Data Operations and Support \(22% of the exam content\)](#)
- [Domain 4: Security and Governance \(16% of the exam content\)](#)
- [In-Scope AWS Services](#)

Introduction

The AWS Certified Data Engineer - Associate (DEA-C01) exam validates a candidate's ability to implement data pipelines and to monitor, troubleshoot, and optimize cost and performance issues in accordance with best practices.

The exam also validates a candidate's ability to complete the following tasks:

- Ingest and transform data, and orchestrate data pipelines while applying programming concepts.
- Choose an optimal data store, design data models, catalog data schemas, and manage data lifecycles.
- Operationalize, maintain, and monitor data pipelines. Analyze data and ensure data quality.
- Implement appropriate authentication, authorization, data encryption, privacy, and governance. Enable logging.

Target Candidate Description

The target candidate should have the equivalent of 2–3 years of experience in data engineering. The target candidate should understand the effects of volume, variety, and velocity on data ingestion, transformation, modeling, security, governance, privacy, schema design, and optimal data store design. Additionally, the target candidate should have at least 1–2 years of hands-on experience with AWS services.

The target candidate should have the following general IT knowledge:

- Setup and maintenance of extract, transform, and load (ETL) pipelines from ingestion to destination
- Application of high-level but language-agnostic programming concepts as required by the pipeline
- How to use Git commands for source control
- How to use data lakes to store data
- General concepts for networking, storage, and compute

The target candidate should have the following AWS knowledge:

- How to use AWS services to accomplish the tasks listed in the Introduction section of this exam guide
- An understanding of the AWS services for encryption, governance, protection, and logging of all data that is part of data pipelines
- The ability to compare AWS services to understand the cost, performance, and functional differences between services
- How to structure SQL queries and how to run SQL queries on AWS services
- An understanding of how to analyze data, verify data quality, and ensure data consistency by using AWS services

The following list contains job tasks that the target candidate is not expected to be able to perform. This list is non-exhaustive. These tasks are out of scope for the exam:

- Develop machine learning models
- Develop applications
- Develop business intelligence (BI) dashboards
- Develop data visualization tools
- Develop data science algorithms

Exam Content

The exam consists of multiple-choice and multiple-response questions. Multiple-choice questions have one correct response and three incorrect responses (distractors). Multiple-response questions have two or more correct responses out of five or more options.

Select one or more responses that best complete the statement or answer the question. Distractors, or incorrect answers, are response options that a candidate with incomplete knowledge or skill might choose. Distractors are generally plausible responses that match the content area.

Unanswered questions are scored as incorrect; there is no penalty for guessing.

The exam includes unscored questions that do not affect your score. AWS collects information about performance on these unscored questions to evaluate these questions for future use as scored questions. These unscored questions are not identified on the exam.

The AWS Certified Data Engineer - Associate (DEA-C01) exam has a pass or fail designation. The exam is scored against a minimum standard established by AWS professionals who follow certification industry best practices and guidelines.

Your results for the exam are reported as a scaled score of 100–1,000. The minimum passing score is 750. Your score shows how you performed on the exam as a whole and whether you passed. Scaled scoring models help equate scores across multiple exam forms that might have slightly different difficulty levels.

Content Domains

For detailed information about each domain, see the following sections:

- [the section called “Domain 1: Data Ingestion and Transformation \(36% of the exam content\)”](#)
- [the section called “Domain 2: Data Store Management \(26% of the exam content\)”](#)
- [the section called “Domain 3: Data Operations and Support \(22% of the exam content\)”](#)
- [the section called “Domain 4: Security and Governance \(16% of the exam content\)”](#)

AWS Services for the Exam

The AWS Certified Data Engineer - Associate exam covers specific AWS services that are relevant to data engineers. Understanding which services are in scope can help you focus your preparation efforts.

For detailed information about the AWS services covered in the exam, see the following section:

- [In-Scope AWS Services](#)

Domain 1: Data Ingestion and Transformation (36% of the exam content)

This domain accounts for 36% of the exam content.

Rubriques

- [Task 1.1: Perform data ingestion](#)
- [Task 1.2: Transform and process data](#)
- [Task 1.3: Orchestrate data pipelines](#)

- [Task 1.4: Apply programming concepts](#)

Task 1.1: Perform data ingestion

Knowledge of:

- Throughput and latency characteristics for services that ingest data
- Data ingestion patterns (for example, frequency and data history)
- Streaming data ingestion
- Batch data ingestion (for example, scheduled ingestion, event-driven ingestion)
- Replayability of data ingestion pipelines
- Stateful and stateless data transactions

Skills in:

- Reading data from streaming sources (for example, Amazon Kinesis, Amazon Managed Streaming for Apache Kafka [Amazon MSK], Amazon DynamoDB Streams, Database Migration Service [DMS], Glue, Amazon Redshift)
- Reading data from batch sources (for example, Amazon S3, Glue, Amazon EMR, DMS, Amazon Redshift, Lambda, Amazon AppFlow)
- Implementing appropriate configuration options for batch ingestion
- Consuming data APIs
- Setting up schedulers by using Amazon EventBridge, Apache Airflow, or time-based schedules for jobs and crawlers
- Setting up event triggers (for example, Amazon S3 Event Notifications, EventBridge)
- Calling a Lambda function from Amazon Kinesis
- Creating allowlists for IP addresses to allow connections to data sources
- Implementing throttling and overcoming rate limits (for example, DynamoDB, Amazon RDS, Kinesis)
- Managing fan-in and fan-out for streaming data distribution

Task 1.2: Transform and process data

Knowledge of:

- Creation of ETL pipelines based on business requirements
- Volume, velocity, and variety of data (for example, structured data, unstructured data)
- Cloud computing and distributed computing
- How to use Apache Spark to process data
- Intermediate data staging locations

Skills in:

- Optimizing container usage for performance needs (for example, Amazon Elastic Kubernetes Service [Amazon EKS], Amazon Elastic Container Service [Amazon ECS])
- Connecting to different data sources (for example, Java Database Connectivity [JDBC], Open Database Connectivity [ODBC])
- Integrating data from multiple sources
- Optimizing costs while processing data
- Implementing data transformation services based on requirements (for example, Amazon EMR, Glue, Lambda, Amazon Redshift)
- Transforming data between formats (for example, from .csv to Apache Parquet)
- Troubleshooting and debugging common transformation failures and performance issues
- Creating data APIs to make data available to other systems by using services

Task 1.3: Orchestrate data pipelines

Knowledge of:

- How to integrate various services to create ETL pipelines
- Event-driven architecture
- How to configure services for data pipelines based on schedules or dependencies
- Serverless workflows

Skills in:

- Using orchestration services to build workflows for data ETL pipelines (for example, Lambda, EventBridge, Amazon Managed Workflows for Apache Airflow [Amazon MWAA], Step Functions, Glue workflows)

- Building data pipelines for performance, availability, scalability, resiliency, and fault tolerance
- Implementing and maintaining serverless workflows
- Using notification services to send alerts (for example, Amazon Simple Notification Service [Amazon SNS], Amazon Simple Queue Service [Amazon SQS])

Task 1.4: Apply programming concepts

Knowledge of:

- Continuous integration and continuous delivery (CI/CD) (implementation, testing, and deployment of data pipelines)
- SQL queries (for data source queries and data transformations)
- Infrastructure as code (IaC) for repeatable deployments (for example, Cloud Development Kit [CDK], CloudFormation)
- Distributed computing
- Data structures and algorithms (for example, graph data structures and tree data structures)
- SQL query optimization

Skills in:

- Optimizing code to reduce runtime for data ingestion and transformation
- Configuring Lambda functions to meet concurrency and performance needs
- Performing SQL queries to transform data (for example, Amazon Redshift stored procedures)
- Structuring SQL queries to meet data pipeline requirements
- Using Git commands to perform actions such as creating, updating, cloning, and branching repositories
- Using the Serverless Application Model (SAM) to package and deploy serverless data pipelines (for example, Lambda functions, Step Functions, DynamoDB tables)
- Using and mounting storage volumes from within Lambda functions

Domain 2: Data Store Management (26% of the exam content)

This domain accounts for 26% of the exam content.

Rubriques

- [Task 2.1: Choose a data store](#)
- [Task 2.2: Understand data cataloging systems](#)
- [Task 2.3: Manage the lifecycle of data](#)
- [Task 2.4: Design data models and schema evolution](#)

Task 2.1: Choose a data store

Knowledge of:

- Storage platforms and their characteristics
- Storage services and configurations for specific performance demands
- Data storage formats (for example, .csv, .txt, Parquet)
- How to align data storage with data migration requirements
- How to determine the appropriate storage solution for specific access patterns
- How to manage locks to prevent access to data (for example, Amazon Redshift, Amazon RDS)

Skills in:

- Implementing the appropriate storage services for specific cost and performance requirements (for example, Amazon Redshift, Amazon EMR, Lake Formation, Amazon RDS, DynamoDB, Amazon Kinesis Data Streams, Amazon MSK)
- Configuring the appropriate storage services for specific access patterns and requirements (for example, Amazon Redshift, Amazon EMR, Lake Formation, Amazon RDS, DynamoDB)
- Applying storage services to appropriate use cases (for example, Amazon S3)
- Integrating migration tools into data processing systems (for example, Transfer Family)
- Implementing data migration or remote access methods (for example, Amazon Redshift federated queries, Amazon Redshift materialized views, Amazon Redshift Spectrum)

Task 2.2: Understand data cataloging systems

Knowledge of:

- How to create a data catalog

- Data classification based on requirements
- Components of metadata and data catalogs

Skills in:

- Using data catalogs to consume data from the data's source
- Building and referencing a data catalog (for example, Glue Data Catalog, Apache Hive metastore)
- Discovering schemas and using Glue crawlers to populate data catalogs
- Synchronizing partitions with a data catalog
- Creating new source or target connections for cataloging (for example, Glue)

Task 2.3: Manage the lifecycle of data

Knowledge of:

- Appropriate storage solutions to address hot and cold data requirements
- How to optimize the cost of storage based on the data lifecycle
- How to delete data to meet business and legal requirements
- Data retention policies and archiving strategies
- How to protect data with appropriate resiliency and availability

Skills in:

- Performing load and unload operations to move data between Amazon S3 and Amazon Redshift
- Managing S3 Lifecycle policies to change the storage tier of S3 data
- Expiring data when it reaches a specific age by using S3 Lifecycle policies
- Managing S3 versioning and DynamoDB TTL

Task 2.4: Design data models and schema evolution

Knowledge of:

- Data modeling concepts
- How to ensure accuracy and trustworthiness of data by using data lineage

- Best practices for indexing, partitioning strategies, compression, and other data optimization techniques
- How to model structured, semi-structured, and unstructured data
- Schema evolution techniques

Skills in:

- Designing schemas for Amazon Redshift, DynamoDB, and Lake Formation
- Addressing changes to the characteristics of data
- Performing schema conversion (for example, by using the Schema Conversion Tool [SCT] and DMS Schema Conversion)
- Establishing data lineage by using tools (for example, Amazon SageMaker ML Lineage Tracking)

Domain 3: Data Operations and Support (22% of the exam content)

This domain accounts for 22% of the exam content.

Rubriques

- [Task 3.1: Automate data processing by using services](#)
- [Task 3.2: Analyze data by using services](#)
- [Task 3.3: Maintain and monitor data pipelines](#)
- [Task 3.4: Ensure data quality](#)

Task 3.1: Automate data processing by using services

Knowledge of:

- How to maintain and troubleshoot data processing for repeatable business outcomes
- API calls for data processing
- Which services accept scripting (for example, Amazon EMR, Amazon Redshift, Glue)

Skills in:

- Orchestrating data pipelines (for example, Amazon MWAA, Step Functions)

- Troubleshooting Amazon managed workflows
- Calling SDKs to access Amazon features from code
- Using the features of services to process data (for example, Amazon EMR, Amazon Redshift, Glue)
- Consuming and maintaining data APIs
- Preparing data transformation (for example, Glue DataBrew)
- Querying data (for example, Amazon Athena)
- Using Lambda to automate data processing
- Managing events and schedulers (for example, EventBridge)

Task 3.2: Analyze data by using services

Knowledge of:

- Tradeoffs between provisioned services and serverless services
- SQL queries (for example, SELECT statements with multiple qualifiers or JOIN clauses)
- How to visualize data for analysis
- When and how to apply cleansing techniques
- Data aggregation, rolling average, grouping, and pivoting

Skills in:

- Visualizing data by using services and tools (for example, Glue DataBrew, Amazon QuickSight)
- Verifying and cleaning data (for example, Lambda, Athena, QuickSight, Jupyter Notebooks, Amazon SageMaker Data Wrangler)
- Using Athena to query data or to create views
- Using Athena notebooks that use Apache Spark to explore data

Task 3.3: Maintain and monitor data pipelines

Knowledge of:

- How to log application data
- Best practices for performance tuning

- How to log access to services
- Amazon Macie, CloudTrail, and Amazon CloudWatch

Skills in:

- Extracting logs for audits
- Deploying logging and monitoring solutions to facilitate auditing and traceability
- Using notifications during monitoring to send alerts
- Troubleshooting performance issues
- Using CloudTrail to track API calls
- Troubleshooting and maintaining pipelines (for example, Glue, Amazon EMR)
- Using Amazon CloudWatch Logs to log application data (with a focus on configuration and automation)
- Analyzing logs with services (for example, Athena, Amazon EMR, Amazon OpenSearch Service, CloudWatch Logs Insights, big data application logs)

Task 3.4: Ensure data quality

Knowledge of:

- Data sampling techniques
- How to implement data skew mechanisms
- Data validation (data completeness, consistency, accuracy, and integrity)
- Data profiling

Skills in:

- Running data quality checks while processing the data (for example, checking for empty fields)
- Defining data quality rules (for example, Glue DataBrew)
- Investigating data consistency (for example, Glue DataBrew)

Domain 4: Security and Governance (16% of the exam content)

This domain accounts for 16% of the exam content.

Rubriques

- [Task 4.1: Apply authentication mechanisms](#)
- [Task 4.2: Apply authorization mechanisms](#)
- [Task 4.3: Ensure data encryption and masking](#)
- [Task 4.4: Prepare logs for audit](#)
- [Task 4.5: Understand data privacy and governance](#)

Task 4.1: Apply authentication mechanisms

Knowledge of:

- VPC security networking concepts
- Differences between managed services and unmanaged services
- Authentication methods (password-based, certificate-based, and role-based)
- Differences between managed policies and customer managed policies

Skills in:

- Updating VPC security groups
- Creating and updating IAM groups, roles, endpoints, and services
- Creating and rotating credentials for password management (for example, Secrets Manager)
- Setting up IAM roles for access (for example, Lambda, Amazon API Gateway, CLI, CloudFormation)
- Applying IAM policies to roles, endpoints, and services (for example, S3 Access Points, PrivateLink)

Task 4.2: Apply authorization mechanisms

Knowledge of:

- Authorization methods (role-based, policy-based, tag-based, and attribute-based)
- Principle of least privilege as it applies to security
- Role-based access control and expected access patterns

- Methods to protect data from unauthorized access across services

Skills in:

- Creating custom IAM policies when a managed policy does not meet the needs
- Storing application and database credentials (for example, Secrets Manager, Systems Manager Parameter Store)
- Providing database users, groups, and roles access and authority in a database (for example, for Amazon Redshift)
- Managing permissions through Lake Formation (for Amazon Redshift, Amazon EMR, Athena, and Amazon S3)

Task 4.3: Ensure data encryption and masking

Knowledge of:

- Data encryption options available in analytics services (for example, Amazon Redshift, Amazon EMR, Glue)
- Differences between client-side encryption and server-side encryption
- Protection of sensitive data
- Data anonymization, masking, and key salting

Skills in:

- Applying data masking and anonymization according to compliance laws or company policies
- Using encryption keys to encrypt or decrypt data (for example, Key Management Service [KMS])
- Configuring encryption across account boundaries
- Enabling encryption in transit for data

Task 4.4: Prepare logs for audit

Knowledge of:

- How to log application data
- How to log access to services

- Centralized logs

Skills in:

- Using CloudTrail to track API calls
- Using CloudWatch Logs to store application logs
- Using CloudTrail Lake for centralized logging queries
- Analyzing logs by using services (for example, Athena, CloudWatch Logs Insights, Amazon OpenSearch Service)
- Integrating various services to perform logging (for example, Amazon EMR in cases of large volumes of log data)

Task 4.5: Understand data privacy and governance

Knowledge of:

- How to protect personally identifiable information (PII)
- Data sovereignty

Skills in:

- Granting permissions for data sharing (for example, data sharing for Amazon Redshift)
- Implementing PII identification (for example, Macie with Lake Formation)
- Implementing data privacy strategies to prevent backups or replications of data to disallowed Regions
- Managing configuration changes that have occurred in an account (for example, Config)

In-Scope AWS Services

The following list contains AWS services and features that are in scope for the AWS Certified Data Engineer - Associate (DEA-C01) exam. This list is non-exhaustive and is subject to change. AWS offerings appear in categories that align with the offerings' primary functions.

Rubriques

- [Analytics](#)

- [Application Integration](#)
- [Compute](#)
- [Containers](#)
- [Database](#)
- [Developer Tools](#)
- [Management and Governance](#)
- [Migration and Transfer](#)
- [Networking and Content Delivery](#)
- [Security, Identity, and Compliance](#)
- [Storage](#)

Analytics

- Amazon Athena
- Amazon EMR
- AWS Glue
- AWS Glue DataBrew
- AWS Lake Formation
- Amazon Kinesis Data Firehose
- Amazon Kinesis Data Streams
- Amazon Managed Service for Apache Flink
- Amazon Managed Streaming for Apache Kafka (Amazon MSK)
- Amazon OpenSearch Service
- Amazon QuickSight
- Amazon Redshift

Application Integration

- Amazon AppFlow
- Amazon EventBridge
- Amazon Managed Workflows for Apache Airflow (Amazon MWAA)

- Amazon Simple Notification Service (Amazon SNS)
- Amazon Simple Queue Service (Amazon SQS)
- AWS Step Functions

Compute

- Amazon EC2
- AWS Lambda

Containers

- Amazon Elastic Container Registry (Amazon ECR)
- Amazon Elastic Container Service (Amazon ECS)
- Amazon Elastic Kubernetes Service (Amazon EKS)

Database

- Amazon Aurora
- Amazon DynamoDB
- Amazon ElastiCache
- Amazon RDS
- Amazon Redshift

Developer Tools

- AWS Cloud Development Kit (AWS CDK)
- AWS CloudFormation
- AWS CodeBuild
- AWS CodeCommit
- AWS CodeDeploy
- AWS CodePipeline
- AWS Serverless Application Model (AWS SAM)

Management and Governance

- AWS CloudTrail
- Amazon CloudWatch
- AWS Config
- AWS Systems Manager

Migration and Transfer

- AWS Database Migration Service (AWS DMS)
- AWS Schema Conversion Tool (AWS SCT)
- AWS Transfer Family

Networking and Content Delivery

- Amazon API Gateway
- Amazon CloudFront
- AWS PrivateLink
- Amazon Route 53
- Amazon VPC

Security, Identity, and Compliance

- AWS Identity and Access Management (IAM)
- AWS Key Management Service (AWS KMS)
- Amazon Macie
- AWS Secrets Manager

Storage

- Amazon Elastic Block Store (Amazon EBS)
- Amazon Elastic File System (Amazon EFS)
- Amazon S3

- Amazon S3 Glacier

Professional Exams

AWS Professional level certification exams validate advanced technical skills and experience in designing and implementing AWS solutions.

Rubriques

- [AWS Certified Solutions Architect - Professional \(SAP-C02\)](#)
- [AWS Certified DevOps Engineer - Professional \(DOP-C02\)](#)

AWS Certified Solutions Architect - Professional (SAP-C02)

The AWS Certified Solutions Architect - Professional (SAP-C02) exam is intended for individuals who perform a solutions architect role. The exam validates a candidate's advanced technical skills and experience in designing optimized AWS solutions that are based on the AWS Well-Architected Framework.

Rubriques

- [Introduction](#)
- [Target Candidate Description](#)
- [Exam Content](#)
- [Content Domains](#)
- [AWS Services for the Exam](#)
- [Domain 1: Design Solutions for Organizational Complexity \(26% of the exam content\)](#)
- [Domain 2: Design for New Solutions \(29% of the exam content\)](#)
- [Domain 3: Migration Planning \(15% of the exam content\)](#)
- [Domain 4: Cost Control \(10% of the exam content\)](#)
- [Domain 5: Continuous Improvement for Existing Solutions \(20% of the exam content\)](#)
- [In-Scope AWS Services](#)
- [Out-of-Scope AWS Services](#)

Introduction

The AWS Certified Solutions Architect - Professional (SAP-C02) exam is intended for individuals who perform a solutions architect role. The exam validates a candidate's advanced technical skills and experience in designing optimized AWS solutions that are based on the AWS Well-Architected Framework.

The exam also validates a candidate's ability to complete the following tasks within the scope of the AWS Well-Architected Framework:

- Design for organizational complexity.
- Design for new solutions.
- Continuously improve existing solutions.
- Accelerate workload migration and modernization.

Target Candidate Description

The target candidate has 2 or more years of experience in using AWS services to design and implement cloud solutions. This candidate has the ability to evaluate cloud application requirements and make architectural recommendations for deployment of applications on AWS. This candidate also can provide expert guidance about architectural design that extends across multiple applications and projects within a complex organization.

The following list contains job tasks that the target candidate is not expected to be able to perform and knowledge that the target candidate is not expected to have. This list is non-exhaustive. These tasks and knowledge are out of scope for the exam:

- Frontend development for mobile apps
- 12-factor app methodology
- In-depth knowledge of operating systems

Exam Content

The exam contains one or more of the following question types:

- Multiple choice: Has one correct response and three incorrect responses (distractors).

- **Multiple response:** Has two or more correct responses out of five or more response options. You must select all the correct responses to receive credit for the question.

Unanswered questions are scored as incorrect. There is no penalty for guessing. The exam includes 65 questions that affect your score.

The exam includes unscored questions that do not affect your score. AWS collects information about performance on these unscored questions to evaluate these questions for future use as scored questions. These unscored questions are not identified on the exam.

The AWS Certified Solutions Architect - Professional (SAP-C02) exam has a pass or fail designation. The exam is scored against a minimum standard established by AWS professionals who follow certification industry best practices and guidelines.

Your results for the exam are reported as a scaled score of 100–1,000. The minimum passing score is 750. Your score shows how you performed on the exam as a whole and whether you passed. Scaled scoring models help equate scores across multiple exam forms that might have slightly different difficulty levels.

Content Domains

For detailed information about each domain, see the following sections:

- [the section called “Domain 1: Design Solutions for Organizational Complexity \(26% of the exam content\)”](#)
- [the section called “Domain 2: Design for New Solutions \(29% of the exam content\)”](#)
- [the section called “Domain 3: Migration Planning \(15% of the exam content\)”](#)
- [the section called “Domain 4: Cost Control \(10% of the exam content\)”](#)
- [the section called “Domain 5: Continuous Improvement for Existing Solutions \(20% of the exam content\)”](#)

AWS Services for the Exam

The AWS Certified Solutions Architect - Professional exam covers specific AWS services that are relevant to solutions architects. Understanding which services are in scope and which are out of scope can help you focus your preparation efforts.

For detailed information about the AWS services covered in the exam, see the following sections:

- [In-Scope AWS Services](#)
- [Out-of-Scope AWS Services](#)

Domain 1: Design Solutions for Organizational Complexity (26% of the exam content)

This domain accounts for 26% of the exam content.

Rubriques

- [Task 1.1: Architect network connectivity strategies](#)
- [Task 1.2: Prescribe security controls](#)
- [Task 1.3: Design reliable and resilient architectures](#)
- [Task 1.4: Design a multi-account environment](#)
- [Task 1.5: Determine cost optimization and visibility strategies](#)

Task 1.1: Architect network connectivity strategies

Knowledge of:

- Global Infrastructure
- networking concepts (for example, Amazon Virtual Private Cloud [Amazon VPC], Direct Connect, VPN, transitive routing, container services)
- Hybrid DNS concepts (for example, Amazon Route 53 Resolver, on-premises DNS integration)
- Network segmentation (for example, subnetting, IP addressing, connectivity among VPCs)
- Network traffic monitoring

Skills in:

- Evaluating connectivity options for multiple VPCs
- Evaluating connectivity options for on-premises, co-location, and cloud integration
- Selecting Regions and Availability Zones based on network and latency requirements
- Troubleshooting traffic flows by using tools

- Using service endpoints for service integrations

Task 1.2: Prescribe security controls

Knowledge of:

- Identity and Access Management (IAM) and IAM Identity Center
- Route tables, security groups, and network ACLs
- Encryption keys and certificate management (for example, Key Management Service [KMS], Certificate Manager [ACM])
- security, identity, and compliance tools (for example, CloudTrail, Identity and Access Management Access Analyzer, Security Hub, Amazon Inspector)

Skills in:

- Evaluating cross-account access management
- Integrating with third-party identity providers
- Deploying encryption strategies for data at rest and data in transit
- Developing a strategy for centralized security event notifications and auditing

Task 1.3: Design reliable and resilient architectures

Knowledge of:

- Recovery time objectives (RTOs) and recovery point objectives (RPOs)
- Disaster recovery strategies (for example, using Elastic Disaster Recovery, pilot light, warm standby, and multi-site)
- Data backup and restoration

Skills in:

- Designing disaster recovery solutions based on RTO and RPO requirements
- Implementing architectures to automatically recover from failure
- Developing the optimal architecture by considering scale-up and scale-out options
- Designing an effective backup and restoration strategy

Task 1.4: Design a multi-account environment

Knowledge of:

- Organizations and Control Tower
- Multi-account event notifications
- resource sharing across environments

Skills in:

- Evaluating the most appropriate account structure for organizational requirements
- Recommending a strategy for central logging and event notifications
- Developing a multi-account governance model

Task 1.5: Determine cost optimization and visibility strategies

Knowledge of:

- cost and usage monitoring tools (for example, Trusted Advisor, Pricing Calculator, Cost Explorer, Budgets)
- purchasing options (for example, Reserved Instances, Savings Plans, Spot Instances)
- rightsizing visibility tools (for example, Compute Optimizer, Amazon Simple Storage Service [Amazon S3] Storage Lens)

Skills in:

- Monitoring cost and usage with tools
- Developing an effective tagging strategy that maps costs to business units
- Understanding how purchasing options affect cost and performance

Domain 2: Design for New Solutions (29% of the exam content)

This domain accounts for 29% of the exam content.

Rubriques

- [Task 2.1: Design a deployment strategy to meet business requirements](#)

- [Task 2.2: Design a solution to ensure business continuity](#)
- [Task 2.3: Design for application security](#)
- [Task 2.4: Design secure application tiers](#)
- [Task 2.5: Select appropriate storage solutions](#)

Task 2.1: Design a deployment strategy to meet business requirements

Knowledge of:

- Infrastructure as code (IaC) (for example, CloudFormation)
- Continuous integration and continuous delivery (CI/CD)
- Change management processes
- Configuration management tools (for example, Systems Manager)

Skills in:

- Determining an application or upgrade path for new services and features
- Selecting services to develop deployment strategies and implement appropriate rollback mechanisms
- Adopting managed services as needed to reduce infrastructure provisioning and patching overhead
- Making advanced technologies accessible by delegating complex development and deployment tasks to

Task 2.2: Design a solution to ensure business continuity

Knowledge of:

- Global Infrastructure
- networking concepts (for example, Route 53, routing methods)
- Disaster recovery strategies and methodologies (for example, backup and restore, pilot light, warm standby, active-active)
- Disaster recovery scenarios (for example, Availability Zone failure, Region failure)

Skills in:

- Implementing the appropriate disaster recovery strategy to meet business requirements
- Implementing architectures to automatically recover from failure
- Implementing multi-Region architectures to support business continuity
- Designing self-healing architectures

Task 2.3: Design for application security

Knowledge of:

- shared responsibility model
- Identity and Access Management (IAM) and IAM Identity Center
- security, identity, and compliance tools (for example, Security Hub, Amazon Inspector, Amazon GuardDuty, Config)
- Encryption and key management (for example, Key Management Service [KMS], Certificate Manager [ACM])

Skills in:

- Designing a strategy for managing multiple accounts
- Designing a strategy for secure application connectivity to services
- Designing a strategy for encryption key management
- Designing a strategy for secure network boundaries

Task 2.4: Design secure application tiers

Knowledge of:

- Security groups and network ACLs
- Load balancing concepts (for example, Application Load Balancer, Network Load Balancer)
- WAF
- service endpoints

Skills in:

- Designing network security controls (for example, security groups, network ACLs)

- Designing web application security using services (for example, WAF, CloudFront)
- Designing secure VPC architectures (for example, private subnets, public subnets)
- Designing secure access to services (for example, service endpoints, VPC endpoints)

Task 2.5: Select appropriate storage solutions

Knowledge of:

- storage services and their features (for example, Amazon S3, Amazon EBS, Amazon EFS, Amazon FSx)
- Storage access patterns
- Storage performance metrics
- Storage tiering and lifecycle management

Skills in:

- Selecting the appropriate storage service based on access patterns
- Designing storage solutions for different access patterns
- Optimizing storage performance
- Designing storage lifecycle management strategies

Domain 3: Migration Planning (15% of the exam content)

This domain accounts for 15% of the exam content.

Rubriques

- [Task 3.1: Select existing workloads and processes for potential migration](#)
- [Task 3.2: Determine the optimal migration strategy for existing workloads](#)
- [Task 3.3: Determine a migration strategy for existing on-premises data](#)

Task 3.1: Select existing workloads and processes for potential migration

Knowledge of:

- Application discovery and assessment

- Application portfolio assessment
- Total cost of ownership (TCO) calculation
- Application migration tools (for example, Application Discovery Service, Migration Hub)

Skills in:

- Evaluating cloud readiness of applications
- Identifying application dependencies
- Calculating TCO for applications
- Prioritizing workloads for migration

Task 3.2: Determine the optimal migration strategy for existing workloads

Knowledge of:

- The 6 Rs of migration (rehost, replatform, repurchase, refactor, retire, retain)
- Application migration tools and services (for example, Application Migration Service, Database Migration Service)
- Data transfer services (for example, DataSync, Transfer Family, Snow Family)
- networking services for migration (for example, Direct Connect, VPN)

Skills in:

- Selecting the appropriate migration strategy for applications
- Selecting the appropriate database migration tools and services
- Selecting the appropriate data transfer services
- Designing network architectures for migration

Task 3.3: Determine a migration strategy for existing on-premises data

Knowledge of:

- Data migration tools and services (for example, Database Migration Service, DataSync)
- Data transfer services (for example, Snow Family)

- Database migration strategies (for example, homogeneous, heterogeneous)
- Data validation and reconciliation

Skills in:

- Selecting the appropriate data migration tools and services
- Designing data migration strategies
- Implementing data validation and reconciliation processes
- Designing strategies for large-scale data migrations

Domain 4: Cost Control (10% of the exam content)

This domain accounts for 10% of the exam content.

Rubriques

- [Task 4.1: Design cost-optimized storage solutions](#)
- [Task 4.2: Design cost-optimized compute solutions](#)
- [Task 4.3: Design cost-optimized database solutions](#)
- [Task 4.4: Design cost-optimized network architectures](#)

Task 4.1: Design cost-optimized storage solutions

Knowledge of:

- storage services and their cost models (for example, Amazon S3, Amazon EBS, Amazon EFS, Amazon FSx)
- Storage tiering and lifecycle management
- Storage access patterns and their impact on cost
- Data transfer costs

Skills in:

- Selecting the most cost-effective storage service for specific workloads
- Implementing storage tiering and lifecycle policies to optimize costs

- Optimizing data transfer costs
- Implementing cost-effective backup and archival strategies

Task 4.2: Design cost-optimized compute solutions

Knowledge of:

- compute services and their cost models (for example, Amazon EC2, Lambda, Amazon ECS, Amazon EKS)
- purchasing options (for example, On-Demand Instances, Reserved Instances, Savings Plans, Spot Instances)
- Compute scaling strategies
- Compute rightsizing

Skills in:

- Selecting the most cost-effective compute service for specific workloads
- Implementing appropriate purchasing options to optimize costs
- Implementing auto scaling to match capacity with demand
- Implementing compute rightsizing strategies

Task 4.3: Design cost-optimized database solutions

Knowledge of:

- database services and their cost models (for example, Amazon RDS, Amazon DynamoDB, Amazon Redshift)
- Database instance sizing
- Database scaling options (for example, read replicas, sharding)
- Database purchasing options (for example, Reserved Instances)

Skills in:

- Selecting the most cost-effective database service for specific workloads
- Implementing appropriate database scaling strategies

- Implementing database purchasing options to optimize costs
- Implementing database rightsizing strategies

Task 4.4: Design cost-optimized network architectures

Knowledge of:

- networking services and their cost models (for example, VPC, Direct Connect, VPN, NAT Gateway)
- Data transfer costs
- Content delivery and edge services (for example, CloudFront)
- Network topology design

Skills in:

- Designing network architectures to minimize data transfer costs
- Implementing content delivery strategies to reduce data transfer costs
- Selecting the most cost-effective connectivity options
- Optimizing NAT Gateway costs

Domain 5: Continuous Improvement for Existing Solutions (20% of the exam content)

This domain accounts for 20% of the exam content.

Rubriques

- [Task 5.1: Troubleshoot solution architectures](#)
- [Task 5.2: Improve solution architectures for reliability](#)
- [Task 5.3: Improve solution architectures for performance](#)
- [Task 5.4: Improve operational excellence](#)

Task 5.1: Troubleshoot solution architectures

Knowledge of:

- monitoring and logging services (for example, Amazon CloudWatch, CloudTrail, X-Ray)
- Troubleshooting methodologies
- Root cause analysis
- service limits and quotas

Skills in:

- Analyzing logs and metrics to identify issues
- Implementing monitoring and alerting strategies
- Performing root cause analysis
- Implementing remediation strategies

Task 5.2: Improve solution architectures for reliability

Knowledge of:

- Well-Architected Framework (Reliability pillar)
- High availability and fault tolerance
- Disaster recovery strategies
- Service quotas and limits

Skills in:

- Implementing multi-AZ and multi-Region architectures
- Implementing auto scaling and self-healing architectures
- Implementing service quotas and limits management
- Implementing disaster recovery strategies

Task 5.3: Improve solution architectures for performance

Knowledge of:

- Well-Architected Framework (Performance Efficiency pillar)
- Performance monitoring and optimization

- Caching strategies
- Database performance optimization

Skills in:

- Implementing caching strategies (for example, Amazon ElastiCache, CloudFront)
- Optimizing database performance
- Implementing performance monitoring and alerting
- Selecting appropriate instance types and sizes

Task 5.4: Improve operational excellence

Knowledge of:

- Well-Architected Framework (Operational Excellence pillar)
- Infrastructure as code (IaC)
- CI/CD pipelines
- Monitoring and observability

Skills in:

- Implementing infrastructure as code (for example, CloudFormation, CDK)
- Implementing CI/CD pipelines (for example, CodePipeline, CodeBuild, CodeDeploy)
- Implementing monitoring and observability (for example, CloudWatch, X-Ray)
- Implementing automated remediation

In-Scope AWS Services

The following list contains AWS services and features that are in scope for the AWS Certified Solutions Architect - Professional (SAP-C02) exam. This list is non-exhaustive and is subject to change. AWS offerings appear in categories that align with the offerings' primary functions.

Rubriques

- [Analytics](#)

- [Application Integration](#)
- [Business Applications](#)
- [Compute](#)
- [Containers](#)
- [Database](#)
- [Developer Tools](#)
- [Frontend Web and Mobile](#)
- [Management and Governance](#)
- [Migration and Transfer](#)
- [Networking and Content Delivery](#)
- [Security, Identity, and Compliance](#)
- [Storage](#)

Analytics

- Amazon Athena
- Amazon EMR
- AWS Glue
- Amazon Kinesis
- Amazon Managed Streaming for Apache Kafka (Amazon MSK)
- Amazon OpenSearch Service
- Amazon QuickSight
- Amazon Redshift

Application Integration

- Amazon AppFlow
- AWS AppSync
- Amazon EventBridge
- Amazon MQ

- Amazon Simple Notification Service (Amazon SNS)
- Amazon Simple Queue Service (Amazon SQS)
- AWS Step Functions

Business Applications

- Amazon Simple Email Service (Amazon SES)

Compute

- AWS Batch
- Amazon EC2
- Amazon EC2 Auto Scaling
- AWS Elastic Beanstalk
- AWS Lambda
- AWS Outposts
- AWS Serverless Application Repository
- AWS Wavelength

Containers

- Amazon Elastic Container Registry (Amazon ECR)
- Amazon Elastic Container Service (Amazon ECS)
- Amazon Elastic Kubernetes Service (Amazon EKS)
- AWS Fargate

Database

- Amazon Aurora
- Amazon DocumentDB (with MongoDB compatibility)
- Amazon DynamoDB
- Amazon ElastiCache

- Amazon Keyspaces (for Apache Cassandra)
- Amazon Neptune
- Amazon RDS
- Amazon Redshift

Developer Tools

- AWS Cloud Development Kit (AWS CDK)
- AWS CloudShell
- AWS CodeArtifact
- AWS CodeBuild
- AWS CodeCommit
- AWS CodeDeploy
- AWS CodePipeline
- AWS CodeStar
- AWS X-Ray

Frontend Web and Mobile

- AWS Amplify
- AWS AppSync
- AWS Device Farm

Management and Governance

- AWS Auto Scaling
- AWS Backup
- AWS CloudFormation
- AWS CloudTrail
- Amazon CloudWatch
- AWS Command Line Interface (AWS CLI)

- AWS Compute Optimizer
- AWS Config
- AWS Control Tower
- AWS License Manager
- AWS Management Console
- AWS Organizations
- AWS Resource Access Manager
- AWS Service Catalog
- AWS Systems Manager
- AWS Trusted Advisor
- AWS Well-Architected Tool

Migration and Transfer

- AWS Application Discovery Service
- AWS Application Migration Service
- AWS Database Migration Service (AWS DMS)
- AWS DataSync
- AWS Migration Hub
- AWS Schema Conversion Tool (AWS SCT)
- AWS Snow Family
- AWS Transfer Family

Networking and Content Delivery

- Amazon API Gateway
- AWS App Mesh
- AWS Client VPN
- Amazon CloudFront
- AWS Cloud Map

- AWS Direct Connect
- Elastic Load Balancing
- AWS Global Accelerator
- AWS PrivateLink
- Amazon Route 53
- AWS Site-to-Site VPN
- AWS Transit Gateway
- Amazon VPC
- AWS VPN

Security, Identity, and Compliance

- AWS Certificate Manager (ACM)
- AWS CloudHSM
- Amazon Cognito
- Amazon Detective
- AWS Directory Service
- AWS Firewall Manager
- Amazon GuardDuty
- AWS Identity and Access Management (IAM)
- Amazon Inspector
- AWS Key Management Service (AWS KMS)
- Amazon Macie
- AWS Network Firewall
- AWS Resource Access Manager (AWS RAM)
- AWS Secrets Manager
- AWS Security Hub
- AWS Shield
- AWS Single Sign-On

- AWS WAF

Storage

- AWS Backup
- Amazon Elastic Block Store (Amazon EBS)
- Amazon Elastic File System (Amazon EFS)
- Amazon FSx (for all types)
- Amazon S3
- Amazon S3 Glacier
- AWS Storage Gateway

Out-of-Scope AWS Services

The following list contains AWS services and features that are out of scope for the AWS Certified Solutions Architect - Professional (SAP-C02) exam. This list is non-exhaustive and is subject to change.

Rubriques

- [Business Applications](#)
- [Compute](#)
- [Game Tech](#)
- [Internet of Things \(IoT\)](#)
- [Machine Learning](#)
- [Media Services](#)
- [Robotics](#)
- [Satellite](#)

Business Applications

- Alexa for Business
- Amazon Chime

- Amazon Connect
- Amazon WorkDocs
- Amazon WorkMail

Compute

- Amazon AppStream 2.0
- Amazon Lightsail
- AWS Mainframe Modernization
- Amazon WorkSpaces

Game Tech

- Amazon GameLift
- Amazon Lumberyard

Internet of Things (IoT)

- AWS IoT 1-Click
- AWS IoT Analytics
- AWS IoT Button
- AWS IoT Core
- AWS IoT Device Defender
- AWS IoT Device Management
- AWS IoT Events
- AWS IoT Greengrass
- AWS IoT SiteWise
- AWS IoT Things Graph

Machine Learning

- Amazon Augmented AI

- Amazon CodeGuru
- Amazon Comprehend
- AWS DeepComposer
- AWS DeepLens
- AWS DeepRacer
- Amazon DevOps Guru
- Amazon Forecast
- Amazon Fraud Detector
- Amazon Kendra
- Amazon Lex
- Amazon Lookout for Equipment
- Amazon Lookout for Metrics
- Amazon Lookout for Vision
- Amazon Monitron
- Amazon Personalize
- Amazon Polly
- Amazon Rekognition
- Amazon SageMaker
- Amazon Textract
- Amazon Transcribe
- Amazon Translate

Media Services

- Amazon Elastic Transcoder
- Amazon Interactive Video Service
- Amazon Kinesis Video Streams
- AWS Elemental MediaConnect
- AWS Elemental MediaConvert

- AWS Elemental MediaLive
- AWS Elemental MediaPackage
- AWS Elemental MediaStore
- AWS Elemental MediaTailor

Robotics

- AWS RoboMaker

Satellite

- AWS Ground Station

AWS Certified DevOps Engineer - Professional (DOP-C02)

The AWS Certified DevOps Engineer - Professional (DOP-C02) exam is intended for individuals who perform a DevOps engineer role. The exam validates a candidate's technical expertise in provisioning, operating, and managing distributed systems and services on AWS.

Rubriques

- [Introduction](#)
- [Target Candidate Description](#)
- [Exam Content](#)
- [Content Domains](#)
- [AWS Services for the Exam](#)
- [Domain 1: SDLC Automation \(22% of the exam content\)](#)
- [Domain 2: Configuration Management and IaC \(17% of the exam content\)](#)
- [Domain 3: Resilient Cloud Solutions \(15% of the exam content\)](#)
- [Domain 4: Monitoring and Logging \(15% of the exam content\)](#)
- [Domain 5: Incident and Event Response \(14% of the exam content\)](#)
- [Domain 6: Security and Compliance \(17% of the exam content\)](#)
- [In-Scope AWS Services](#)

Introduction

The AWS Certified DevOps Engineer - Professional (DOP-C02) exam is intended for individuals who perform a DevOps engineer role. The exam validates a candidate's technical expertise in provisioning, operating, and managing distributed systems and services on AWS.

The exam also validates a candidate's ability to complete the following tasks:

- Implement and manage continuous delivery systems and methodologies on AWS.
- Implement and automate security controls, governance processes, and compliance validation.
- Define and deploy monitoring, metrics, and logging systems on AWS.
- Implement systems that are highly available, scalable, and self-healing on AWS.
- Design, manage, and maintain tools to automate operational processes.

Target Candidate Description

The target candidate should have 2 or more years of experience in provisioning, operating, and managing AWS environments. The target candidate also has experience with the software development lifecycle and programming and/or scripting.

The target candidate should have the following experience:

- Experience in building highly automated infrastructure
- Experience in administering operating systems
- Experience with modern development and operations processes and methodologies

The target candidate should have experience in securing AWS infrastructure.

The following list contains job tasks that the target candidate is not expected to be able to perform. This list is non-exhaustive. These tasks are out of scope for the exam:

- Possess advanced networking knowledge (for example, advanced routing algorithms, failover techniques).
- Provide deep-level security recommendations to developers.
- Design, query, and optimize the performance of databases.

- Develop full-stack application code.

Exam Content

The exam consists of multiple-choice and multiple-response questions. Multiple-choice questions have one correct response and three incorrect responses (distractors). Multiple-response questions have two or more correct responses out of five or more options.

Select one or more responses that best complete the statement or answer the question. Distractors, or incorrect answers, are response options that a candidate with incomplete knowledge or skill might choose. Distractors are generally plausible responses that match the content area.

Unanswered questions are scored as incorrect; there is no penalty for guessing.

The exam includes unscored questions that do not affect your score. AWS collects information about performance on these unscored questions to evaluate these questions for future use as scored questions. These unscored questions are not identified on the exam.

Content Domains

For detailed information about each domain, see the following sections:

- [the section called “Domain 1: SDLC Automation \(22% of the exam content\)”](#)
- [the section called “Domain 2: Configuration Management and IaC \(17% of the exam content\)”](#)
- [the section called “Domain 3: Resilient Cloud Solutions \(15% of the exam content\)”](#)
- [the section called “Domain 4: Monitoring and Logging \(15% of the exam content\)”](#)
- [the section called “Domain 5: Incident and Event Response \(14% of the exam content\)”](#)
- [the section called “Domain 6: Security and Compliance \(17% of the exam content\)”](#)

AWS Services for the Exam

The AWS Certified DevOps Engineer - Professional exam covers specific AWS services that are relevant to DevOps engineers. Understanding which services are in scope can help you focus your preparation efforts.

For detailed information about the AWS services covered in the exam, see the following section:

- [In-Scope AWS Services](#)

Domain 1: SDLC Automation (22% of the exam conten)

This domain accounts for 22% of the exam content.

Rubriques

- [Task 1.1: Implement CI/CD pipelines](#)
- [Task 1.2: Integrate automated testing into CI/CD pipelines](#)
- [Task 1.3: Build and manage artifacts](#)
- [Task 1.4: Implement deployment strategies for instance, container, and serverless environments](#)

Task 1.1: Implement CI/CD pipelines

Knowledge of:

- Software development lifecycle (SDLC) concepts, phases, and models
- Pipeline deployment patterns for single- and multi-account environments

Skills in:

- Configuring code, image, and artifact repositories
- Using version control to integrate pipelines with application environments
- Setting up build processes (for example, CodeBuild)
- Managing build and deployment secrets (for example, Secrets Manager, Systems Manager Parameter Store)
- Determining appropriate deployment strategies (for example, CodeDeploy)

Task 1.2: Integrate automated testing into CI/CD pipelines

Knowledge of:

- Different types of tests (for example, unit tests, integration tests, acceptance tests, user interface tests, security scans)
- Reasonable use of different types of tests at different stages of the CI/CD pipeline

Skills in:

- Running builds or tests when generating pull requests or code merges (for example, CodeBuild)
- Running load/stress tests, performance benchmarking, and application testing at scale
- Measuring application health based on application exit codes
- Automating unit tests and code coverage
- Invoking services in a pipeline for testing

Task 1.3: Build and manage artifacts

Knowledge of:

- Artifact use cases and secure management
- Methods to create and generate artifacts
- Artifact lifecycle considerations

Skills in:

- Creating and configuring artifact repositories (for example, CodeArtifact, Amazon S3, Amazon Elastic Container Registry [Amazon ECR])
- Configuring build tools for generating artifacts (for example, CodeBuild, Lambda)
- Automating Amazon EC2 instance and container image build processes (for example, EC2 Image Builder)

Task 1.4: Implement deployment strategies for instance, container, and serverless environments

Knowledge of:

- Deployment methodologies for various platforms (for example, Amazon EC2, Amazon Elastic Container Service [Amazon ECS], Amazon Elastic Kubernetes Service [Amazon EKS], Lambda)
- Application storage patterns (for example, Amazon Elastic File System [Amazon EFS], Amazon S3, Amazon Elastic Block Store [Amazon EBS])
- Mutable deployment patterns in contrast to immutable deployment patterns
- Tools and services available for distributing code (for example, CodeDeploy, EC2 Image Builder)

Skills in:

- Configuring security permissions to allow access to artifact repositories (for example, Identity and Access Management [IAM], CodeArtifact)
- Configuring deployment agents (for example, CodeDeploy agent)
- Troubleshooting deployment issues
- Using different deployment methods (for example, blue/green, canary)

Domain 2: Configuration Management and IaC (17% of the exam content)

This domain accounts for 17% of the exam content.

Rubriques

- [Task 2.1: Define cloud infrastructure and reusable components to provision and manage systems throughout their lifecycle](#)
- [Task 2.2: Automate configuration management to enforce desired state](#)
- [Task 2.3: Implement compliance as code](#)

Task 2.1: Define cloud infrastructure and reusable components to provision and manage systems throughout their lifecycle

Knowledge of:

- Infrastructure as code (IaC) options and tools for
- Change management processes for IaC-based platforms
- Configuration management services and strategies

Skills in:

- Composing and deploying IaC templates (for example, Serverless Application Model [SAM], CloudFormation, Cloud Development Kit [CDK])
- Applying CloudFormation StackSets across multiple accounts and Regions
- Determining optimal configuration management services (for example, Systems Manager, AppConfig)

- Implementing infrastructure and application configuration management
- Implementing IaC lifecycle management (for example, drift detection, stack updates)
- Implementing IaC security controls (for example, CloudFormation Guard)

Task 2.2: Automate configuration management to enforce desired state

Knowledge of:

- Configuration management tools and services
- Desired state management
- Immutable infrastructure concepts

Skills in:

- Implementing configuration management services (for example, Systems Manager State Manager, Config)
- Implementing patch management (for example, Systems Manager Patch Manager)
- Implementing application configuration management (for example, AppConfig)
- Implementing desired state management (for example, Systems Manager State Manager)
- Implementing infrastructure and application configuration management

Task 2.3: Implement compliance as code

Knowledge of:

- Compliance as code concepts
- Compliance and security scanning tools
- Compliance and security remediation strategies

Skills in:

- Implementing compliance scanning tools (for example, Config, Security Hub, Amazon Inspector)
- Implementing automated remediation (for example, Config, Systems Manager)
- Implementing security scanning tools (for example, Amazon Inspector, Security Hub)

- Implementing infrastructure security scanning (for example, cfn-nag, CloudFormation Guard)

Domain 3: Resilient Cloud Solutions (15% of the exam content)

This domain accounts for 15% of the exam content.

Rubriques

- [Task 3.1: Implement highly available solutions to meet resilience and business requirements](#)
- [Task 3.2: Implement self-healing and scaling solutions to meet resilience and business requirements](#)
- [Task 3.3: Implement backup and restore strategies](#)

Task 3.1: Implement highly available solutions to meet resilience and business requirements

Knowledge of:

- Multi-AZ and multi-Region deployments (for example, compute layer, data layer)
- SLAs
- Replication and failover methods for stateful services
- Techniques to achieve high availability (for example, Multi-AZ, multi-Region)

Skills in:

- Implementing high availability for compute (for example, Auto Scaling groups, Spot Instance interruption handling)
- Implementing high availability for storage (for example, Amazon S3 replication, Amazon RDS Multi-AZ)
- Implementing high availability for databases (for example, Amazon RDS read replicas, Amazon Aurora Global Database)
- Implementing high availability for networking (for example, Route 53 routing policies, Global Accelerator)
- Implementing high availability for containers (for example, Amazon ECS, Amazon EKS)
- Implementing high availability for serverless (for example, Lambda)

Task 3.2: Implement self-healing and scaling solutions to meet resilience and business requirements

Knowledge of:

- Auto scaling concepts and features
- Self-healing architectures
- Scaling strategies (for example, horizontal, vertical)
- Immutable infrastructure concepts

Skills in:

- Implementing auto scaling for compute (for example, EC2 Auto Scaling, Application Auto Scaling)
- Implementing auto scaling for containers (for example, Amazon ECS, Amazon EKS)
- Implementing auto scaling for serverless (for example, Lambda provisioned concurrency)
- Implementing self-healing architectures (for example, health checks, automated recovery)
- Implementing immutable infrastructure (for example, blue/green deployments, canary deployments)

Task 3.3: Implement backup and restore strategies

Knowledge of:

- Backup and restore strategies
- Disaster recovery strategies (for example, backup and restore, pilot light, warm standby, active-active)
- RPO and RTO concepts

Skills in:

- Implementing backup and restore for compute (for example, AMIs, EC2 instance recovery)
- Implementing backup and restore for storage (for example, Amazon S3 versioning, Backup)
- Implementing backup and restore for databases (for example, Amazon RDS snapshots, DynamoDB point-in-time recovery)
- Implementing disaster recovery strategies (for example, CloudFormation, Backup)
- Implementing cross-Region and cross-account backup strategies

Domain 4: Monitoring and Logging (15% of the exam content)

This domain accounts for 15% of the exam content.

Rubriques

- [Task 4.1: Configure the collection, aggregation, and storage of logs and metrics](#)
- [Task 4.2: Audit, monitor, and visualize data to gain operational insights](#)
- [Task 4.3: Implement automated responses to events and alerts](#)

Task 4.1: Configure the collection, aggregation, and storage of logs and metrics

Knowledge of:

- Log collection and aggregation services (for example, Amazon CloudWatch Logs, CloudTrail)
- Metric collection and aggregation services (for example, CloudWatch, Amazon Managed Service for Prometheus)
- Log and metric storage options (for example, CloudWatch Logs, S3, Amazon OpenSearch Service)
- Log and metric retention strategies

Skills in:

- Configuring log collection for services (for example, CloudWatch Logs, CloudTrail)
- Configuring metric collection for services (for example, CloudWatch, Amazon Managed Service for Prometheus)
- Configuring log aggregation (for example, CloudWatch Logs subscription filters, CloudWatch Logs Insights)
- Configuring metric aggregation (for example, CloudWatch metric math, CloudWatch metric streams)
- Implementing log and metric storage solutions (for example, S3 lifecycle policies, CloudWatch Logs retention)

Task 4.2: Audit, monitor, and visualize data to gain operational insights

Knowledge of:

- Monitoring and observability concepts
- Visualization tools and services (for example, CloudWatch dashboards, Amazon Managed Grafana)
- Auditing services and features (for example, CloudTrail, Config)
- Log and metric analysis techniques

Skills in:

- Creating and configuring dashboards (for example, CloudWatch dashboards, Amazon Managed Grafana)
- Analyzing logs (for example, CloudWatch Logs Insights, Amazon OpenSearch Service)
- Analyzing metrics (for example, CloudWatch metric math, CloudWatch anomaly detection)
- Implementing auditing solutions (for example, CloudTrail, Config)
- Creating custom metrics (for example, CloudWatch embedded metric format, CloudWatch custom metrics)

Task 4.3: Implement automated responses to events and alerts

Knowledge of:

- Event-driven architectures
- Alerting services and features (for example, CloudWatch alarms, Amazon SNS)
- Automated response services and features (for example, CloudWatch Events, Amazon EventBridge, Lambda)
- Incident management processes

Skills in:

- Creating and configuring alarms (for example, CloudWatch alarms)
- Implementing automated responses to events (for example, EventBridge rules, Lambda functions)
- Implementing automated responses to alarms (for example, CloudWatch alarm actions)
- Implementing notification systems (for example, SNS, Amazon SQS)
- Implementing automated remediation (for example, Systems Manager Automation, Config remediation)

Domain 5: Incident and Event Response (14% of the exam content)

This domain accounts for 14% of the exam content.

Rubriques

- [Task 5.1: Implement automated alerting and event management](#)
- [Task 5.2: Implement automated healing and recovery](#)
- [Task 5.3: Troubleshoot system and application failures](#)

Task 5.1: Implement automated alerting and event management

Knowledge of:

- Event-driven architectures
- Alerting services and features (for example, CloudWatch alarms, Amazon SNS)
- Event management services and features (for example, EventBridge, CloudWatch Events)
- Incident management processes

Skills in:

- Creating and configuring alarms (for example, CloudWatch alarms)
- Implementing event management (for example, EventBridge rules, CloudWatch Events)
- Implementing notification systems (for example, SNS, Amazon SQS)
- Implementing incident management processes (for example, OpsCenter, Incident Manager)
- Implementing automated escalation (for example, SNS, Incident Manager)

Task 5.2: Implement automated healing and recovery

Knowledge of:

- Self-healing architectures
- Automated recovery services and features (for example, Auto Scaling, Amazon RDS Multi-AZ)
- Automated remediation services and features (for example, Systems Manager Automation, Config remediation)
- Health check mechanisms

Skills in:

- Implementing self-healing architectures (for example, Auto Scaling, health checks)
- Implementing automated recovery (for example, RDS Multi-AZ, Route 53 health checks)
- Implementing automated remediation (for example, Systems Manager Automation, Config remediation)
- Implementing health checks (for example, Route 53 health checks, load balancer health checks)
- Implementing circuit breakers and fallbacks

Task 5.3: Troubleshoot system and application failures

Knowledge of:

- Troubleshooting methodologies
- Common system and application failure modes
- Diagnostic tools and services (for example, CloudWatch Logs Insights, X-Ray)
- Root cause analysis techniques

Skills in:

- Analyzing logs and metrics to identify issues (for example, CloudWatch Logs Insights, X-Ray)
- Troubleshooting system failures (for example, EC2 instances, RDS databases)
- Troubleshooting application failures (for example, Lambda functions, containers)
- Performing root cause analysis
- Implementing corrective actions

Domain 6: Security and Compliance (17% of the exam content)

This domain accounts for 17% of the exam content.

Rubriques

- [Task 6.1: Implement application security](#)
- [Task 6.2: Implement infrastructure security](#)
- [Task 6.3: Automate security and compliance validation](#)

Task 6.1: Implement application security

Knowledge of:

- Application security best practices
- Authentication and authorization services and features (for example, IAM, Amazon Cognito)
- Secrets management services and features (for example, Secrets Manager, Systems Manager Parameter Store)
- Data protection services and features (for example, KMS, Certificate Manager)

Skills in:

- Implementing authentication and authorization (for example, IAM roles, Amazon Cognito)
- Implementing secrets management (for example, Secrets Manager, Systems Manager Parameter Store)
- Implementing data protection (for example, encryption at rest, encryption in transit)
- Implementing secure API endpoints (for example, API Gateway with WAF)
- Implementing security headers and content security policies

Task 6.2: Implement infrastructure security

Knowledge of:

- Infrastructure security best practices
- Network security services and features (for example, security groups, network ACLs, WAF)
- Identity and access management services and features (for example, IAM, Organizations)
- Threat detection services and features (for example, Amazon GuardDuty, Amazon Inspector)

Skills in:

- Implementing network security (for example, security groups, network ACLs)
- Implementing identity and access management (for example, IAM policies, service control policies)
- Implementing threat detection (for example, GuardDuty, Inspector)
- Implementing infrastructure protection (for example, Shield, WAF)

- Implementing secure VPC architectures

Task 6.3: Automate security and compliance validation

Knowledge of:

- Compliance frameworks and requirements
- Security assessment services and features (for example, Config, Security Hub)
- Automated remediation services and features (for example, Config remediation, Systems Manager Automation)
- Security testing methodologies

Skills in:

- Implementing compliance validation (for example, Config rules, Security Hub standards)
- Implementing automated remediation (for example, Config remediation, Systems Manager Automation)
- Implementing security testing (for example, penetration testing, vulnerability scanning)
- Implementing security monitoring (for example, CloudTrail, CloudWatch Logs)
- Implementing security automation (for example, Security Hub, Config)

In-Scope AWS Services

The following list contains AWS services and features that are in scope for the AWS Certified DevOps Engineer - Professional (DOP-C02) exam. This list is non-exhaustive and is subject to change. AWS offerings appear in categories that align with the offerings' primary functions.

Rubriques

- [Analytics](#)
- [Application Integration](#)
- [Compute](#)
- [Containers](#)
- [Database](#)
- [Developer Tools](#)

- [Management and Governance](#)
- [Networking and Content Delivery](#)
- [Security, Identity, and Compliance](#)
- [Storage](#)

Analytics

- Amazon Athena
- Amazon Kinesis
- Amazon OpenSearch Service

Application Integration

- Amazon EventBridge
- Amazon Simple Notification Service (Amazon SNS)
- Amazon Simple Queue Service (Amazon SQS)
- AWS Step Functions

Compute

- AWS Batch
- Amazon EC2
- Amazon EC2 Auto Scaling
- AWS Elastic Beanstalk
- AWS Lambda
- AWS Serverless Application Model (AWS SAM)

Containers

- Amazon Elastic Container Registry (Amazon ECR)
- Amazon Elastic Container Service (Amazon ECS)
- Amazon Elastic Kubernetes Service (Amazon EKS)
- AWS Fargate

Database

- Amazon Aurora
- Amazon DynamoDB
- Amazon ElastiCache
- Amazon RDS

Developer Tools

- AWS Cloud Development Kit (AWS CDK)
- AWS CloudShell
- AWS CodeArtifact
- AWS CodeBuild
- AWS CodeCommit
- AWS CodeDeploy
- AWS CodePipeline
- AWS CodeStar
- AWS X-Ray

Management and Governance

- AWS Auto Scaling
- AWS Backup
- AWS CloudFormation
- AWS CloudTrail
- Amazon CloudWatch
- AWS Command Line Interface (AWS CLI)
- AWS Compute Optimizer
- AWS Config
- AWS Control Tower
- AWS License Manager
- AWS Management Console

- AWS Organizations
- AWS Resource Access Manager
- AWS Service Catalog
- AWS Systems Manager
- AWS Trusted Advisor

Networking and Content Delivery

- Amazon API Gateway
- Amazon CloudFront
- Elastic Load Balancing
- Amazon Route 53
- Amazon VPC

Security, Identity, and Compliance

- AWS Certificate Manager (ACM)
- Amazon Cognito
- AWS Identity and Access Management (IAM)
- AWS Key Management Service (AWS KMS)
- AWS Secrets Manager
- AWS Security Hub
- AWS Shield
- AWS WAF

Storage

- Amazon Elastic Block Store (Amazon EBS)
- Amazon Elastic File System (Amazon EFS)
- Amazon S3

Specialty Exams

AWS Specialty certification exams validate technical skills and experience in specific technical domains.

Rubriques

- [AWS Certified Machine Learning - Specialty \(MLS-C01\)](#)
- [AWS Certified Security - Specialty \(SCS-C02\)](#)
- [AWS Certified Advanced Networking - Specialty \(ANS-C01\)](#)

AWS Certified Machine Learning - Specialty (MLS-C01)

The AWS Certified Machine Learning - Specialty (MLS-C01) exam is intended for individuals who perform an artificial intelligence and machine learning (AI/ML) development or data science role. The exam validates a candidate's ability to design, build, deploy, optimize, train, tune, and maintain ML solutions for given business problems by using the AWS Cloud.

Rubriques

- [Introduction](#)
- [Target Candidate Description](#)
- [Exam Content](#)
- [Content Domains](#)
- [AWS Services for the Exam](#)
- [Domain 1: Data Engineering \(20% of the exam content\)](#)
- [Domain 2: Exploratory Data Analysis \(24% of the exam content\)](#)
- [Domain 3: Modeling \(36% of the exam content\)](#)
- [Domain 4: Machine Learning Implementation and Operations \(20% of the exam content\)](#)
- [In-Scope AWS Services](#)
- [Out-of-Scope AWS Services](#)

Introduction

The AWS Certified Machine Learning - Specialty (MLS-C01) exam is intended for individuals who perform an artificial intelligence and machine learning (AI/ML) development or data science role. The exam validates a candidate's ability to design, build, deploy, optimize, train, tune, and maintain ML solutions for given business problems by using the AWS Cloud.

The exam also validates a candidate's ability to complete the following tasks:

- Select and justify the appropriate ML approach for a given business problem.
- Identify appropriate AWS services to implement ML solutions.
- Design and implement scalable, cost-optimized, reliable, and secure ML solutions.

Target Candidate Description

The target candidate should have 2 or more years of experience developing, architecting, and running ML or deep learning workloads in the AWS Cloud.

The target candidate should have the following AWS knowledge:

- The ability to express the intuition behind basic ML algorithms
- Experience performing basic hyperparameter optimization
- Experience with ML and deep learning frameworks
- The ability to follow model-training best practices
- The ability to follow deployment best practices
- The ability to follow operational best practices

The following list contains knowledge that the target candidate is not expected to have. This list is non-exhaustive. Knowledge in the following areas is out of scope for the exam:

- Extensive or complex algorithm development
- Extensive hyperparameter optimization
- Complex mathematical proofs and computations
- Advanced networking and network design
- Advanced database, security, and DevOps concepts

- DevOps-related tasks for Amazon EMR

Exam Content

There are two types of questions on the exam:

- Multiple choice: Has one correct response and three incorrect responses (distractors)
- Multiple response: Has two or more correct responses out of five or more response options

Select one or more responses that best complete the statement or answer the question. Distractors, or incorrect answers, are response options that a candidate with incomplete knowledge or skill might choose. Distractors are generally plausible responses that match the content area.

Unanswered questions are scored as incorrect; there is no penalty for guessing. The exam includes 50 questions that affect your score.

The exam includes 15 unscored questions that do not affect your score. AWS collects information about performance on these unscored questions to evaluate these questions for future use as scored questions. These unscored questions are not identified on the exam.

The AWS Certified Machine Learning - Specialty (MLS-C01) exam has a pass or fail designation. The exam is scored against a minimum standard established by AWS professionals who follow certification industry best practices and guidelines.

Your results for the exam are reported as a scaled score of 100–1,000. The minimum passing score is 750. Your score shows how you performed on the exam as a whole and whether you passed. Scaled scoring models help equate scores across multiple exam forms that might have slightly different difficulty levels.

Content Domains

For detailed information about each domain, see the following sections:

- [the section called “Domain 1: Data Engineering \(20% of the exam content\)”](#)
- [the section called “Domain 2: Exploratory Data Analysis \(24% of the exam content\)”](#)
- [the section called “Domain 3: Modeling \(36% of the exam content\)”](#)
- [the section called “Domain 4: Machine Learning Implementation and Operations \(20% of the exam content\)”](#)

AWS Services for the Exam

The AWS Certified Machine Learning - Specialty exam covers specific AWS services that are relevant to machine learning workloads. Understanding which services are in scope and which are out of scope can help you focus your preparation efforts.

For detailed information about the AWS services covered in the exam, see the following sections:

- [In-Scope AWS Services](#)
- [Out-of-Scope AWS Services](#)

Domain 1: Data Engineering (20% of the exam content)

This domain accounts for 20% of the exam content.

Rubriques

- [Task 1.1: Create data repositories for ML](#)
- [Task 1.2: Identify and implement a data ingestion solution](#)
- [Task 1.3: Identify and implement a data transformation solution](#)

Task 1.1: Create data repositories for ML

- Identify data sources (for example, content and location, primary sources such as user data).
- Determine storage mediums (for example, databases, Amazon S3, Amazon Elastic File System [Amazon EFS], Amazon Elastic Block Store [Amazon EBS]).

Task 1.2: Identify and implement a data ingestion solution

- Identify data job styles and job types (for example, batch load, streaming).
- Orchestrate data ingestion pipelines (batch-based ML workloads and streaming-based ML workloads).
 - Amazon Kinesis
 - Amazon Data Firehose
 - Amazon EMR
 - Glue

- Amazon Managed Service for Apache Flink
- Schedule jobs.

Task 1.3: Identify and implement a data transformation solution

- Transform data in transit (ETL, Glue, Amazon EMR, Batch).
- Handle ML-specific data by using MapReduce (for example, Apache Hadoop, Apache Spark, Apache Hive).

Domain 2: Exploratory Data Analysis (24% of the exam content)

This domain accounts for 24% of the exam content.

Rubriques

- [Task 2.1: Sanitize and prepare data for modeling](#)
- [Task 2.2: Perform feature engineering](#)
- [Task 2.3: Analyze and visualize data for ML](#)

Task 2.1: Sanitize and prepare data for modeling

- Identify and handle missing data, corrupt data, and stop words.
- Format, normalize, augment, and scale data.
- Determine whether there is sufficient labeled data.
 - Identify mitigation strategies.
 - Use data labelling tools (for example, Amazon Mechanical Turk).

Task 2.2: Perform feature engineering

- Identify and extract features from datasets, including from data sources such as text, speech, images, and public datasets.
- Analyze and evaluate feature engineering concepts (for example, binning, tokenization, outliers, synthetic features, one-hot encoding, reducing dimensionality of data).

Task 2.3: Analyze and visualize data for ML

- Create graphs (for example, scatter plots, time series, histograms, box plots).
- Interpret descriptive statistics (for example, correlation, summary statistics, p-value).
- Perform cluster analysis (for example, hierarchical, diagnosis, elbow plot, cluster size).

Domain 3: Modeling (36% of the exam content)

This domain accounts for 36% of the exam content.

Rubriques

- [Task 3.1: Frame business problems as ML problems](#)
- [Task 3.2: Select the appropriate model\(s\) for a given ML problem](#)
- [Task 3.3: Train ML models](#)
- [Task 3.4: Perform hyperparameter optimization](#)
- [Task 3.5: Evaluate ML models](#)

Task 3.1: Frame business problems as ML problems

- Determine when to use and when not to use ML.
- Know the difference between supervised and unsupervised learning.
- Select from among classification, regression, forecasting, clustering, recommendation, and foundation models.

Task 3.2: Select the appropriate model(s) for a given ML problem

- XGBoost, logistic regression, k-means, linear regression, decision trees, random forests, RNN, CNN, ensemble, transfer learning, and large language models (LLMs)
- Express the intuition behind models.

Task 3.3: Train ML models

- Split data between training and validation (for example, cross validation).

- Understand optimization techniques for ML training (for example, gradient descent, loss functions, convergence).
- Choose appropriate compute resources (for example GPU or CPU, distributed or non-distributed).
 - Choose appropriate compute platforms (Spark or non-Spark).
- Update and retrain models.
 - Batch or real-time/online

Task 3.4: Perform hyperparameter optimization

- Perform regularization.
 - Dropout
 - L1/L2
- Perform cross-validation.
- Initialize models.
- Understand neural network architecture (layers and nodes), learning rate, and activation functions.
- Understand tree-based models (number of trees, number of levels).
- Understand linear models (learning rate).

Task 3.5: Evaluate ML models

- Avoid overfitting or underfitting.
 - Detect and handle bias and variance.
- Evaluate metrics (for example, area under curve [AUC]-receiver operating characteristics [ROC], accuracy, precision, recall, Root Mean Square Error [RMSE], F1 score).
- Interpret confusion matrices.
- Perform offline and online model evaluation (A/B testing).
- Compare models by using metrics (for example, time to train a model, quality of model, engineering costs).
- Perform cross-validation.

Domain 4: Machine Learning Implementation and Operations (20% of the exam content)

This domain accounts for 20% of the exam content.

Rubriques

- [Task 4.1: Build ML solutions for performance, availability, scalability, resiliency, and fault tolerance](#)
- [Task 4.2: Recommend and implement the appropriate ML services and features for a given problem](#)
- [Task 4.3: Apply basic security practices to ML solutions](#)
- [Task 4.4: Deploy and operationalize ML solutions](#)

Task 4.1: Build ML solutions for performance, availability, scalability, resiliency, and fault tolerance

- Log and monitor environments.
 - CloudTrail and Amazon CloudWatch
 - Build error monitoring solutions.
- Deploy to multiple Regions and multiple Availability Zones.
- Create AMIs and golden images.
- Create Docker containers.
- Deploy Auto Scaling groups.
- Rightsize resources (for example, instances, Provisioned IOPS, volumes).
- Perform load balancing.
- Follow best practices.

Task 4.2: Recommend and implement the appropriate ML services and features for a given problem

- ML on (application services), for example:
 - Amazon Polly
 - Amazon Lex
 - Amazon Transcribe

- Amazon Q
- Understand service quotas.
- Determine when to build custom models and when to use Amazon SageMaker built-in algorithms.
- Understand infrastructure (for example, instance types) and cost considerations.
 - Use Spot Instances to train deep learning models by using Batch.

Task 4.3: Apply basic security practices to ML solutions

- Identity and Access Management (IAM)
- S3 bucket policies
- Security groups
- VPCs
- Encryption and anonymization

Task 4.4: Deploy and operationalize ML solutions

- Expose endpoints and interact with them.
- Understand ML models.
- Perform A/B testing.
- Retrain pipelines.
- Debug and troubleshoot ML models.
 - Detect and mitigate drops in performance.
 - Monitor performance of the model.

In-Scope AWS Services

The following list contains AWS services and features that are in scope for the AWS Certified Machine Learning - Specialty (MLS-C01) exam. This list is non-exhaustive and is subject to change. AWS offerings appear in categories that align with the offerings' primary functions.

Rubriques

- [Analytics](#)
- [Compute](#)

- [Containers](#)
- [Database](#)
- [Internet of Things](#)
- [Machine Learning](#)
- [Management and Governance](#)
- [Networking and Content Delivery](#)
- [Security, Identity, and Compliance](#)
- [Storage](#)

Analytics

- Amazon Athena
- Amazon Data Firehose
- Amazon EMR
- AWS Glue
- Amazon Kinesis
- Amazon Kinesis Data Streams
- AWS Lake Formation
- Amazon Managed Service for Apache Flink
- Amazon OpenSearch Service
- Amazon QuickSight

Compute

- AWS Batch
- Amazon EC2
- AWS Lambda

Containers

- Amazon Elastic Container Registry (Amazon ECR)
- Amazon Elastic Container Service (Amazon ECS)

- Amazon Elastic Kubernetes Service (Amazon EKS)
- AWS Fargate

Database

- Amazon Redshift

Internet of Things

- AWS IoT Greengrass

Machine Learning

- Amazon Bedrock
- Amazon Comprehend
- AWS Deep Learning AMIs (DLAMI)
- Amazon Forecast
- Amazon Fraud Detector
- Amazon Lex
- Amazon Kendra
- Amazon Mechanical Turk
- Amazon Polly
- Amazon Q
- Amazon Rekognition
- Amazon SageMaker
- Amazon Textract
- Amazon Transcribe
- Amazon Translate

Management and Governance

- AWS CloudTrail

- Amazon CloudWatch

Networking and Content Delivery

- Amazon VPC

Security, Identity, and Compliance

- AWS Identity and Access Management (IAM)

Storage

- Amazon Elastic Block Store (Amazon EBS)
- Amazon Elastic File System (Amazon EFS)
- Amazon FSx
- Amazon S3

Out-of-Scope AWS Services

The following list contains AWS services and features that are out of scope for the AWS Certified Machine Learning - Specialty (MLS-C01) exam. This list is non-exhaustive and is subject to change. AWS offerings that are entirely unrelated to the target job roles for the exam are excluded from this list.

Rubriques

- [Analytics](#)
- [Machine Learning](#)

Analytics

- AWS Data Pipeline

Machine Learning

- AWS DeepRacer

- Amazon Machine Learning (Amazon ML)

AWS Certified Security - Specialty (SCS-C02)

The AWS Certified Security - Specialty (SCS-C02) exam is intended for individuals who perform a security role. The exam validates a candidate's ability to effectively demonstrate knowledge about securing AWS products and services.

Rubriques

- [Introduction](#)
- [Target Candidate Description](#)
- [Exam Content](#)
- [Content Domains](#)
- [AWS Services for the Exam](#)
- [Domain 1: Threat Detection and Incident Response \(20% of the exam content\)](#)
- [Domain 2: Security Logging and Monitoring \(18% of the exam content\)](#)
- [Domain 3: Infrastructure Security \(16% of the exam content\)](#)
- [Domain 4: Identity and Access Management \(20% of the exam content\)](#)
- [Domain 5: Data Protection \(18% of the exam content\)](#)
- [Domain 6: Management and Security Governance \(8% of the exam content\)](#)
- [In-Scope AWS Services](#)

Introduction

The AWS Certified Security - Specialty (SCS-C02) exam is intended for individuals who perform a security role. The exam validates a candidate's ability to effectively demonstrate knowledge about securing AWS products and services.

The exam also validates whether a candidate has the following:

- An understanding of specialized data classifications and AWS data protection mechanisms
- An understanding of data-encryption methods and AWS mechanisms to implement them
- An understanding of secure internet protocols and AWS mechanisms to implement them

- A working knowledge of AWS security services and features of services to provide a secure production environment
- Competency from 2 or more years of production deployment experience in using AWS security services and features
- The ability to make tradeoff decisions regarding cost, security, and deployment complexity to meet a set of application requirements
- An understanding of security operations and risks

Target Candidate Description

The target candidate should have the equivalent of 3–5 years of experience in designing and implementing security solutions. Additionally, the target candidate should have a minimum of 2 years of hands-on experience in securing AWS workloads.

The target candidate should have the following knowledge:

- The AWS shared responsibility model and its application
- General knowledge of AWS services and deploying cloud solutions
- Security controls for AWS environments and workloads
- Logging and monitoring strategies
- Vulnerability management and security automation
- Ways to integrate AWS security services with third-party tools
- Disaster recovery controls, including backup strategies
- Cryptography and key management
- Identity access management
- Data retention and lifecycle management
- How to troubleshoot security issues
- Multi-account governance and organizational compliance
- Threat detection and incident response strategies

The following list contains job tasks that the target candidate is not expected to be able to perform. This list is non-exhaustive. These tasks are out of scope for the exam:

- Perform penetration testing

- Perform application code reviews
- Design application architecture
- Troubleshoot application integration issues
- Implement security automation systems by using programming languages

Exam Content

The exam consists of multiple-choice and multiple-response questions. Multiple-choice questions have one correct response and three incorrect responses (distractors). Multiple-response questions have two or more correct responses out of five or more options.

Select one or more responses that best complete the statement or answer the question. Distractors, or incorrect answers, are response options that a candidate with incomplete knowledge or skill might choose. Distractors are generally plausible responses that match the content area.

Unanswered questions are scored as incorrect; there is no penalty for guessing.

The exam includes unscored questions that do not affect your score. AWS collects information about performance on these unscored questions to evaluate these questions for future use as scored questions. These unscored questions are not identified on the exam.

The AWS Certified Security - Specialty (SCS-C02) exam has a pass or fail designation. The exam is scored against a minimum standard established by AWS professionals who follow certification industry best practices and guidelines.

Your results for the exam are reported as a scaled score of 100–1,000. The minimum passing score is 750. Your score shows how you performed on the exam as a whole and whether you passed. Scaled scoring models help equate scores across multiple exam forms that might have slightly different difficulty levels.

Content Domains

For detailed information about each domain, see the following sections:

- [the section called “Domain 1: Threat Detection and Incident Response \(20% of the exam content\)”](#)
- [the section called “Domain 2: Security Logging and Monitoring \(18% of the exam content\)”](#)
- [the section called “Domain 3: Infrastructure Security \(16% of the exam content\)”](#)
- [the section called “Domain 4: Identity and Access Management \(20% of the exam content\)”](#)

- [the section called “Domain 5: Data Protection \(18% of the exam content\)”](#)
- [the section called “Domain 6: Management and Security Governance \(8% of the exam content\)”](#)

AWS Services for the Exam

The AWS Certified Security - Specialty exam covers specific AWS services that are relevant to security professionals. Understanding which services are in scope can help you focus your preparation efforts.

For detailed information about the AWS services covered in the exam, see the following section:

- [In-Scope AWS Services](#)

Domain 1: Threat Detection and Incident Response (20% of the exam content)

This domain accounts for 20% of the exam content.

Rubriques

- [Task 1.1: Design and implement an incident response plan](#)
- [Task 1.2: Detect security threats and anomalies by using services](#)
- [Task 1.3: Respond to security incidents and events](#)

Task 1.1: Design and implement an incident response plan

Knowledge of:

- best practices for incident response
- Cloud incidents
- Roles and responsibilities in the incident response plan
- Security Finding Format (ASFF)

Skills in:

- Implementing credential invalidation and rotation strategies in response to compromises (for example, by using Identity and Access Management [IAM] and Secrets Manager)

- Isolating resources
- Designing and implementing playbooks and runbooks for responses to security incidents
- Deploying security services (for example, Security Hub, Amazon Macie, Amazon GuardDuty, Amazon Inspector, Config, Amazon Detective, Identity and Access Management Access Analyzer)
- Configuring integrations with native services and third-party services (for example, by using Amazon EventBridge and the ASFF)

Task 1.2: Detect security threats and anomalies by using services

Knowledge of:

- managed security services that detect threats
- Anomaly and correlation techniques to join data across services
- Visualizations to identify anomalies
- Strategies to centralize security findings

Skills in:

- Evaluating findings from security services (for example, GuardDuty, Security Hub, Macie, Config, IAM Access Analyzer)
- Searching and correlating security threats across services (for example, by using Detective)
- Performing queries to validate security events (for example, by using Amazon Athena)
- Creating metric filters and dashboards to detect anomalous activity (for example, by using Amazon CloudWatch)

Task 1.3: Respond to security incidents and events

Knowledge of:

- Security Incident Response Guide
- Resource isolation mechanisms
- Techniques for root cause analysis
- Data capture mechanisms
- Log analysis for event validation

Skills in:

- Automating remediation by using services (for example, Lambda, Step Functions, EventBridge, Systems Manager runbooks, Security Hub, Config)
- Responding to compromised resources (for example, by isolating Amazon EC2 instances)
- Investigating and analyzing to conduct root cause analysis (for example, by using Detective)
- Capturing relevant forensics data from a compromised resource (for example, Amazon Elastic Block Store [Amazon EBS] volume snapshots, memory dump)
- Querying logs in Amazon S3 for contextual information related to security events (for example, by using Athena)
- Protecting and preserving forensic artifacts (for example, by using S3 Object Lock, isolated forensic accounts, S3 Lifecycle, and S3 replication)
- Preparing services for incidents and recovering services after incidents

Domain 2: Security Logging and Monitoring (18% of the exam content)

This domain accounts for 18% of the exam content.

Rubriques

- [Task 2.1: Design and implement monitoring and alerting to address security events](#)
- [Task 2.2: Troubleshoot security monitoring and alerting](#)
- [Task 2.3: Design and implement a logging solution](#)
- [Task 2.4: Troubleshoot logging solutions](#)
- [Task 2.5: Design a log analysis solution](#)

Task 2.1: Design and implement monitoring and alerting to address security events

Knowledge of:

- services that monitor events and provide alarms (for example, CloudWatch, EventBridge)
- services that automate alerting (for example, Lambda, Amazon Simple Notification Service [Amazon SNS], Security Hub)
- Tools that monitor metrics and baselines (for example, GuardDuty, Systems Manager)

Skills in:

- Analyzing architectures to identify monitoring requirements and sources of data for security monitoring
- Analyzing environments and workloads to determine monitoring requirements
- Designing environment monitoring and workload monitoring based on business and security requirements
- Setting up automated tools and scripts to perform regular audits (for example, by creating custom insights in Security Hub)
- Defining the metrics and thresholds that generate alerts

Task 2.2: Troubleshoot security monitoring and alerting

Knowledge of:

- Configuration of monitoring services (for example, Security Hub)
- Relevant data that indicates security events

Skills in:

- Analyzing the service functionality, permissions, and configuration of resources after an event that did not provide visibility or alerting
- Analyzing and remediating the configuration of a custom application that is not reporting its statistics
- Evaluating logging and monitoring services for alignment with security requirements

Task 2.3: Design and implement a logging solution

Knowledge of:

- services and features that provide logging capabilities (for example, VPC Flow Logs, DNS logs, CloudTrail, Amazon CloudWatch Logs)
- Attributes of logging capabilities (for example, log levels, type, verbosity)
- Log destinations and lifecycle management (for example, retention period)

Skills in:

- Configuring logging for services and applications
- Identifying logging requirements and sources for log ingestion
- Implementing log storage and lifecycle management according to best practices and organizational requirements

Task 2.4: Troubleshoot logging solutions

Knowledge of:

- Capabilities and use cases of services that provide data sources (for example, log level, type, verbosity, cadence, timeliness, immutability)
- services and features that provide logging capabilities (for example, VPC Flow Logs, DNS logs, CloudTrail, CloudWatch Logs)
- Access permissions that are necessary for logging

Skills in:

- Identifying misconfiguration and determining remediation steps for absent access permissions that are necessary for logging (for example, by managing read/write permissions, S3 bucket permissions, public access, and integrity)
- Determining the cause of missing logs and performing remediation steps

Task 2.5: Design a log analysis solution

Knowledge of:

- Services and tools to analyze captured logs (for example, Athena, CloudWatch Logs filter)
- Log analysis features of services (for example, CloudWatch Logs Insights, CloudTrail Insights, Security Hub insights)
- Log format and components (for example, CloudTrail logs)

Skills in:

- Identifying patterns in logs to indicate anomalies and known threats

- Normalizing, parsing, and correlating logs

Domain 3: Infrastructure Security (16% of the exam content)

This domain accounts for 16% of the exam content.

Rubriques

- [Task 3.1: Design and implement security controls for edge services](#)
- [Task 3.2: Design and implement network security controls](#)
- [Task 3.3: Design and implement security controls for compute workloads](#)
- [Task 3.4: Troubleshoot network security](#)

Task 3.1: Design and implement security controls for edge services

Knowledge of:

- Security features on edge services (for example, WAF, load balancers, Amazon Route 53, Amazon CloudFront, Shield)
- Common attacks, threats, and exploits (for example, Open Web Application Security Project [OWASP] Top 10, DDoS)
- Layered web application architecture

Skills in:

- Defining edge security strategies for common use cases (for example, public website, serverless app, mobile app backend)
- Selecting appropriate edge services based on anticipated threats and attacks (for example, OWASP Top 10, DDoS)
- Selecting appropriate protections based on anticipated vulnerabilities and risks (for example, vulnerable software, applications, libraries)
- Defining layers of defense by combining edge security services (for example, CloudFront with WAF and load balancers)
- Applying restrictions at the edge based on various criteria (for example, geography, geolocation, rate limit)
- Activating logs, metrics, and monitoring around edge services to indicate attacks

Task 3.2: Design and implement network security controls

Knowledge of:

- VPC security mechanisms (for example, security groups, network ACLs, Network Firewall)
- Inter-VPC connectivity (for example, Transit Gateway, VPC endpoints)
- Security telemetry sources (for example, Traffic Mirroring, VPC Flow Logs)
- VPN technology, terminology, and usage
- On-premises connectivity options (for example, VPN, Direct Connect)

Skills in:

- Implementing network segmentation based on security requirements (for example, public subnets, private subnets, sensitive VPCs, on-premises connectivity)
- Designing network controls to permit or prevent network traffic as required (for example, by using security groups, network ACLs, and Network Firewall)
- Designing network flows to keep data off the public internet (for example, by using Transit Gateway, VPC endpoints, and Lambda in VPCs)
- Determining which telemetry sources to monitor based on network design, threats, and attacks (for example, load balancer logs, VPC Flow Logs, Traffic Mirroring)
- Determining redundancy and security workload requirements for communication between on-premises environments and the Cloud (for example, by using VPN, VPN over Direct Connect, and MACsec)
- Identifying and removing unnecessary network access
- Managing network configurations as requirements change (for example, by using Firewall Manager)

Task 3.3: Design and implement security controls for compute workloads

Knowledge of:

- Provisioning and maintenance of EC2 instances (for example, patching, inspecting, creation of snapshots and AMIs, use of EC2 Image Builder)
- IAM instance roles and IAM service roles

- Services that scan for vulnerabilities in compute workloads (for example, Amazon Inspector, Amazon Elastic Container Registry [Amazon ECR])
- Host-based security (for example, firewalls, hardening)

Skills in:

- Creating hardened EC2 AMIs
- Applying instance roles and service roles as appropriate to authorize compute workloads
- Scanning EC2 instances and container images for known vulnerabilities
- Applying patches across a fleet of EC2 instances or container images
- Activating host-based security mechanisms (for example, host-based firewalls)
- Analyzing Amazon Inspector findings and determining appropriate mitigation techniques
- Passing secrets and credentials securely to compute workloads

Task 3.4: Troubleshoot network security

Knowledge of:

- How to analyze reachability (for example, by using VPC Reachability Analyzer and Amazon Inspector)
- Fundamental TCP/IP networking concepts (for example, UDP compared with TCP, ports, Open Systems Interconnection [OSI] model, network operating system utilities)
- How to read relevant log sources (for example, Route 53 logs, WAF logs, VPC Flow Logs)

Skills in:

- Identifying, interpreting, and prioritizing problems in network connectivity (for example, by using Amazon Inspector Network Reachability)
- Determining solutions to produce desired network behavior
- Analyzing log sources to identify problems
- Capturing traffic samples for problem analysis (for example, by using Traffic Mirroring)

Domain 4: Identity and Access Management (20% of the exam content)

This domain accounts for 20% of the exam content.

Rubriques

- [Task 4.1: Design, implement, and troubleshoot authentication for resources](#)
- [Task 4.2: Design, implement, and troubleshoot authorization for resources](#)
- [Task 4.3: Design and implement role-based access control](#)

Task 4.1: Design, implement, and troubleshoot authentication for resources

Knowledge of:

- Methods and services for creating and managing identities (for example, federation, identity providers, IAM Identity Center [Single Sign-On], Amazon Cognito)
- Long-term and temporary credentialing mechanisms
- How to troubleshoot authentication issues (for example, by using CloudTrail, IAM Access Advisor, and IAM policy simulator)

Skills in:

- Establishing identity through an authentication system, based on requirements
- Setting up multi-factor authentication (MFA)
- Determining when to use Security Token Service (STS) to issue temporary credentials
- Implementing identity federation (for example, SAML, OIDC)
- Troubleshooting authentication issues using CloudTrail and other services

Task 4.2: Design, implement, and troubleshoot authorization for resources

Knowledge of:

- Different IAM policies (for example, managed policies, inline policies, identity-based policies, resource-based policies, session control policies)
- Components and impact of a policy (for example, Principal, Action, Resource, Condition)

- How to troubleshoot authorization issues (for example, by using CloudTrail, IAM Access Advisor, and IAM policy simulator)
- Policy evaluation logic and policy types
- Permission boundaries and service control policies (SCPs)

Skills in:

- Implementing least privilege access
- Writing and analyzing IAM policies
- Using IAM Access Analyzer to identify unintended access
- Implementing permission boundaries
- Implementing service control policies (SCPs)
- Troubleshooting authorization issues using CloudTrail, IAM Access Advisor, and IAM policy simulator

Task 4.3: Design and implement role-based access control

Knowledge of:

- IAM roles and their use cases
- Cross-account access mechanisms
- Organizations and organizational units (OUs)
- Control Tower and account management

Skills in:

- Designing and implementing IAM roles for different use cases
- Implementing cross-account access
- Designing and implementing Organizations structure
- Implementing service control policies (SCPs) across an organization
- Using Control Tower to manage accounts and implement guardrails

Domain 5: Data Protection (18% of the exam content)

This domain accounts for 18% of the exam content.

Rubriques

- [Task 5.1: Design and implement controls that provide confidentiality and integrity for data in transit](#)
- [Task 5.2: Design and implement controls that provide confidentiality and integrity for data at rest](#)
- [Task 5.3: Design and implement controls to manage the lifecycle of data at rest](#)
- [Task 5.4: Design and implement controls to protect credentials, secrets, and cryptographic key materials](#)

Task 5.1: Design and implement controls that provide confidentiality and integrity for data in transit

Knowledge of:

- TLS concepts
- VPN concepts (for example, IPsec)
- Secure remote access methods (for example, SSH, RDP over Systems Manager Session Manager)
- Systems Manager Session Manager concepts
- How TLS certificates work with various network services and resources (for example, CloudFront, load balancers)

Skills in:

- Designing secure connectivity between and on-premises networks (for example, by using Direct Connect and VPN gateways)
- Designing mechanisms to require encryption when connecting to resources (for example, Amazon RDS, Amazon Redshift, CloudFront, Amazon S3, Amazon DynamoDB, load balancers, Amazon Elastic File System [Amazon EFS], Amazon API Gateway)
- Requiring TLS for API calls (for example, with Amazon S3)
- Designing mechanisms to forward traffic over secure connections (for example, by using Systems Manager and EC2 Instance Connect)

- Designing cross-Region networking by using private VIFs and public VIFs

Task 5.2: Design and implement controls that provide confidentiality and integrity for data at rest

Knowledge of:

- Encryption technique selection (for example, client-side, server-side, symmetric, asymmetric)
- Integrity-checking techniques (for example, hashing algorithms, digital signatures)
- Resource policies (for example, for DynamoDB, Amazon S3, and Key Management Service [KMS])
- IAM roles and policies

Skills in:

- Designing resource policies to restrict access to authorized users (for example, S3 bucket policies, DynamoDB policies)
- Designing mechanisms to prevent unauthorized public access (for example, S3 Block Public Access, prevention of public snapshots and public AMIs)
- Configuring services to activate encryption of data at rest (for example, Amazon S3, Amazon RDS, DynamoDB, Amazon Simple Queue Service [Amazon SQS], Amazon EBS, Amazon EFS)
- Designing mechanisms to protect data integrity by preventing modifications (for example, by using S3 Object Lock, KMS key policies, S3 Glacier Vault Lock, and Backup Vault Lock)
- Designing encryption at rest by using CloudHSM for relational databases (for example, Amazon RDS, RDS Custom, databases on EC2 instances)
- Choosing encryption techniques based on business requirements

Task 5.3: Design and implement controls to manage the lifecycle of data at rest

Knowledge of:

- Lifecycle policies
- Data retention standards

Skills in:

- Designing S3 Lifecycle mechanisms to retain data for required retention periods (for example, S3 Object Lock, S3 Glacier Vault Lock, S3 Lifecycle policy)
- Designing automatic lifecycle management for services and resources (for example, Amazon S3, EBS volume snapshots, RDS volume snapshots, AMIs, container images, CloudWatch log groups, Amazon Data Lifecycle Manager)
- Establishing schedules and retention for Backup across services

Task 5.4: Design and implement controls to protect credentials, secrets, and cryptographic key materials

Knowledge of:

- Secrets Manager
- Systems Manager Parameter Store
- Usage and management of symmetric keys and asymmetric keys (for example, KMS)

Skills in:

- Designing management and rotation of secrets for workloads (for example, database access credentials, API keys, IAM access keys, KMS customer managed keys)
- Designing KMS key policies to limit key usage to authorized users
- Establishing mechanisms to import and remove customer-provided key material
- Implementing secure storage and retrieval of secrets
- Implementing automatic rotation of secrets

Domain 6: Management and Security Governance (8% of the exam content)

This domain accounts for 8% of the exam content.

Rubriques

- [Task 6.1: Develop a strategy to centrally deploy and manage accounts](#)
- [Task 6.2: Implement a secure and consistent deployment strategy for cloud resources](#)
- [Task 6.3: Evaluate the compliance of resources](#)
- [Task 6.4: Identify security gaps through architectural reviews and cost analysis](#)

Task 6.1: Develop a strategy to centrally deploy and manage accounts

Knowledge of:

- Multi-account strategies
- Managed services that allow delegated administration
- Policy-defined guardrails
- Root account best practices
- Cross-account roles

Skills in:

- Deploying and configuring Organizations
- Determining when and how to deploy Control Tower (for example, which services must be deactivated for successful deployment)
- Implementing SCPs as a technical solution to enforce a policy (for example, limitations on the use of a root account, implementation of controls in Control Tower)
- Centrally managing security services and aggregating findings (for example, by using delegated administration and Config aggregators)
- Securing account root user credentials

Task 6.2: Implement a secure and consistent deployment strategy for cloud resources

Knowledge of:

- Deployment best practices with infrastructure as code (IaC) (for example, CloudFormation template hardening and drift detection)
- Best practices for tagging
- Centralized management, deployment, and versioning of services
- Visibility and control over infrastructure

Skills in:

- Using CloudFormation to deploy cloud resources consistently and securely

- Implementing and enforcing multi-account tagging strategies
- Configuring and deploying portfolios of approved services (for example, by using Service Catalog)
- Organizing resources into different groups for management
- Deploying Firewall Manager to enforce policies
- Securely sharing resources across accounts (for example, by using Resource Access Manager [RAM])

Task 6.3: Evaluate the compliance of resources

Knowledge of:

- Data classification by using services
- How to assess, audit, and evaluate the configurations of resources (for example, by using Config)

Skills in:

- Identifying sensitive data by using Macie
- Creating Config rules for detection of noncompliant resources
- Collecting and organizing evidence by using Security Hub and Audit Manager

Task 6.4: Identify security gaps through architectural reviews and cost analysis

Knowledge of:

- cost and usage for anomaly identification
- Strategies to reduce attack surfaces
- Well-Architected Framework

Skills in:

- Identifying anomalies based on resource utilization and trends
- Identifying unused resources by using services and tools (for example, Trusted Advisor, Cost Explorer)
- Using the Well-Architected Tool to identify security gaps

In-Scope AWS Services

The following list contains AWS services and features that are in scope for the AWS Certified Security - Specialty (SCS-C02) exam. This list is non-exhaustive and is subject to change. AWS offerings appear in categories that align with the offerings' primary functions.

Note: Security affects all AWS services. Many services do not appear in this list because the overall service is out of scope, but the security aspects of the service are in scope. For example, a candidate for this exam would not be asked about the steps to set up replication for an S3 bucket. However, the candidate might be asked about configuring an S3 bucket policy.

Rubriques

- [Management and Governance](#)
- [Security, Identity, and Compliance](#)
- [Networking and Content Delivery](#)
- [Compute](#)
- [Containers](#)
- [Storage](#)
- [Database](#)
- [Analytics](#)
- [Application Integration](#)

Management and Governance

- AWS Audit Manager
- AWS CloudFormation
- AWS CloudTrail
- Amazon CloudWatch
- AWS Config
- AWS Control Tower
- AWS Organizations
- AWS Resource Access Manager (AWS RAM)
- AWS Service Catalog

- AWS Systems Manager
- AWS Trusted Advisor
- AWS Well-Architected Tool

Security, Identity, and Compliance

- AWS Artifact
- AWS Certificate Manager (ACM)
- AWS CloudHSM
- Amazon Cognito
- Amazon Detective
- AWS Directory Service
- AWS Firewall Manager
- Amazon GuardDuty
- AWS IAM Identity Center (AWS Single Sign-On)
- AWS Identity and Access Management (IAM)
- Amazon Inspector
- AWS Key Management Service (AWS KMS)
- Amazon Macie
- AWS Network Firewall
- AWS Resource Access Manager (AWS RAM)
- AWS Secrets Manager
- AWS Security Hub
- AWS Security Token Service (AWS STS)
- AWS Shield
- AWS WAF

Networking and Content Delivery

- Amazon API Gateway
- Amazon CloudFront

- AWS Direct Connect
- Elastic Load Balancing (ELB)
- AWS PrivateLink
- Amazon Route 53
- AWS Transit Gateway
- Amazon VPC
- AWS VPN

Compute

- Amazon EC2
- AWS Lambda

Containers

- Amazon Elastic Container Registry (Amazon ECR)
- Amazon Elastic Container Service (Amazon ECS)
- Amazon Elastic Kubernetes Service (Amazon EKS)

Storage

- AWS Backup
- Amazon Elastic Block Store (Amazon EBS)
- Amazon Elastic File System (Amazon EFS)
- Amazon S3
- Amazon S3 Glacier

Database

- Amazon Aurora
- Amazon DynamoDB
- Amazon RDS
- Amazon Redshift

Analytics

- Amazon Athena
- AWS Glue
- Amazon Kinesis
- Amazon OpenSearch Service

Application Integration

- Amazon EventBridge
- Amazon Simple Notification Service (Amazon SNS)
- Amazon Simple Queue Service (Amazon SQS)
- AWS Step Functions

AWS Certified Advanced Networking - Specialty (ANS-C01)

The AWS Certified Advanced Networking - Specialty (ANS-C01) exam is intended for individuals who perform an AWS networking specialist's role. The exam validates a candidate's ability to design, implement, manage, and secure AWS and hybrid network architectures at scale.

Rubriques

- [Introduction](#)
- [Target Candidate Description](#)
- [Exam Content](#)
- [Content Domains](#)
- [AWS Services for the Exam](#)
- [Domain 1: Network Design \(30% of the exam content\)](#)
- [Domain 2: Network Implementation \(26% of the exam content\)](#)
- [Domain 3: Network Management and Operation \(20% of the exam content\)](#)
- [Domain 4: Network Security, Compliance, and Governance \(24% of the exam content\)](#)
- [In-Scope AWS Services](#)
- [Out-of-Scope AWS Services](#)

Introduction

The AWS Certified Advanced Networking - Specialty (ANS-C01) exam is intended for individuals who perform an AWS networking specialist's role. The exam validates a candidate's ability to design, implement, manage, and secure AWS and hybrid network architectures at scale.

The exam also validates a candidate's ability to complete the following tasks:

- Design and develop hybrid and cloud-based networking solutions by using AWS
- Implement core AWS networking services according to AWS best practices
- Operate and maintain hybrid and cloud-based network architecture for all AWS services
- Use tools to deploy and automate hybrid and cloud-based AWS networking tasks
- Implement secure AWS networks using AWS native networking constructs and services

Target Candidate Description

The target candidate should have 5 or more years of networking experience with 2 or more years of cloud and hybrid networking experience.

The target candidate should have the following AWS knowledge:

- AWS networking nuances and how they relate to the integration of AWS services
- AWS security best practices
- AWS compute and storage options and their underlying consistency models

Exam Content

There are two types of questions on the exam:

- Multiple choice: Has one correct response and three incorrect responses (distractors)
- Multiple response: Has two or more correct responses out of five or more response options

Select one or more responses that best complete the statement or answer the question. Distractors, or incorrect answers, are response options that a candidate with incomplete knowledge or skill might choose. Distractors are generally plausible responses that match the content area.

Unanswered questions are scored as incorrect; there is no penalty for guessing.

The exam includes unscored questions that do not affect your score. AWS collects information about performance on these unscored questions to evaluate these questions for future use as scored questions. These unscored questions are not identified on the exam.

The AWS Certified Advanced Networking - Specialty (ANS-C01) exam has a pass or fail designation. The exam is scored against a minimum standard established by AWS professionals who follow certification industry best practices and guidelines.

Your results for the exam are reported as a scaled score of 100–1,000. The minimum passing score is 750. Your score shows how you performed on the exam as a whole and whether you passed. Scaled scoring models help equate scores across multiple exam forms that might have slightly different difficulty levels.

Content Domains

For detailed information about each domain, see the following sections:

- [the section called “Domain 1: Network Design \(30% of the exam content\)”](#)
- [the section called “Domain 2: Network Implementation \(26% of the exam content\)”](#)
- [the section called “Domain 3: Network Management and Operation \(20% of the exam content\)”](#)
- [the section called “Domain 4: Network Security, Compliance, and Governance \(24% of the exam content\)”](#)

AWS Services for the Exam

The AWS Certified Advanced Networking - Specialty exam covers specific AWS services that are relevant to networking specialists. Understanding which services are in scope and which are out of scope can help you focus your preparation efforts.

For detailed information about the AWS services covered in the exam, see the following sections:

- [In-Scope AWS Services](#)
- [Out-of-Scope AWS Services](#)

Domain 1: Network Design (30% of the exam content)

This domain accounts for 30% of the exam content.

Rubriques

- [Task 1.1: Design a solution that incorporates edge network services to optimize user performance and traffic management for global architectures](#)
- [Task 1.2: Design DNS solutions that meet public, private, and hybrid requirements](#)
- [Task 1.3: Design solutions that integrate load balancing to meet high availability, scalability, and security requirements](#)
- [Task 1.4: Define logging and monitoring requirements across and hybrid networks](#)
- [Task 1.5: Design a routing strategy and connectivity architecture between on-premises networks and the Cloud](#)
- [Task 1.6: Design a routing strategy and connectivity architecture that include multiple accounts, Regions, and VPCs to support different connectivity patterns](#)

Task 1.1: Design a solution that incorporates edge network services to optimize user performance and traffic management for global architectures

Knowledge of:

- Design patterns for the usage of content distribution networks (for example, Amazon CloudFront)
- Design patterns for global traffic management (for example, Global Accelerator)
- Integration patterns for content distribution networks and global traffic management with other services (for example, Elastic Load Balancing [ELB], Amazon API Gateway)

Skills in:

- Evaluating requirements of global inbound and outbound traffic from the internet to design an appropriate content distribution solution

Task 1.2: Design DNS solutions that meet public, private, and hybrid requirements

Knowledge of:

- DNS protocol (for example, DNS records, TTL, DNSSEC, DNS delegation, zones)
- DNS logging and monitoring
- Amazon Route 53 features (for example, alias records, traffic policies, resolvers, health checks)
- Integration of Route 53 with other networking services (for example, Amazon VPC)

- Integration of Route 53 with hybrid, multi-account, and multi-Region options
- Domain registration

Skills in:

- Using Route 53 public hosted zones
- Using Route 53 private hosted zones
- Using Route 53 Resolver endpoints in hybrid and architectures
- Using Route 53 for global traffic management
- Creating and managing domain registrations

Task 1.3: Design solutions that integrate load balancing to meet high availability, scalability, and security requirements

Knowledge of:

- How load balancing works at layer 3, layer 4, and layer 7 of the OSI model
- Different types of load balancers and how they meet requirements for network design, high availability, and security
- Connectivity patterns that apply to load balancing based on the use case (for example, internal load balancers, external load balancers)
- Scaling factors for load balancers
- Integrations of load balancers and other services (for example, Global Accelerator, CloudFront, WAF, Route 53, Amazon Elastic Kubernetes Service [Amazon EKS], Certificate Manager [ACM])
- Configuration options for load balancers (for example, proxy protocol, cross-zone load balancing, session affinity [sticky sessions], routing algorithms)
- Configuration options for load balancer target groups (for example, TCP, GENEVE, IP compared with instance)
- Load Balancer Controller for Kubernetes clusters
- Considerations for encryption and authentication with load balancers (for example, TLS termination, TLS passthrough)

Skills in:

- Selecting an appropriate load balancer based on the use case
- Integrating auto scaling with load balancing solutions
- Integrating load balancers with existing application deployments

Task 1.4: Define logging and monitoring requirements across and hybrid networks

Knowledge of:

- Amazon CloudWatch metrics, agents, logs, alarms, dashboards, and insights in architectures to provide visibility
- Transit Gateway Network Manager in architectures to provide visibility
- VPC Reachability Analyzer in architectures to provide visibility
- Flow logs and traffic mirroring in architectures to provide visibility
- Access logging (for example, load balancers, CloudFront)

Skills in:

- Identifying the logging and monitoring requirements
- Recommending appropriate metrics to provide visibility of the network status
- Capturing baseline network performance

Task 1.5: Design a routing strategy and connectivity architecture between on-premises networks and the Cloud

Knowledge of:

- Routing fundamentals (for example, dynamic compared with static, BGP)
- Layer 1 and layer 2 concepts for physical interconnects (for example, VLAN, link aggregation group [LAG], optics, jumbo frames)
- Encapsulation and encryption technologies (for example, Generic Routing Encapsulation [GRE], IPsec)
- Resource sharing across accounts
- Overlay networks

Skills in:

- Identifying the requirements for hybrid connectivity
- Designing a redundant hybrid connectivity model with services (for example, Direct Connect, Site-to-Site VPN)
- Designing BGP routing with BGP attributes to influence the traffic flows based on the desired traffic patterns (load sharing, active/passive)
- Designing for integration of a software-defined wide area network (SD-WAN) with (for example, Transit Gateway Connect, overlay networks)

Task 1.6: Design a routing strategy and connectivity architecture that include multiple accounts, Regions, and VPCs to support different connectivity patterns

Knowledge of:

- Different connectivity patterns and use cases (for example, VPC peering, Transit Gateway, PrivateLink)
- Capabilities and advantages of VPC sharing
- IP subnets and solutions accounting for IP address overlaps

Skills in:

- Connecting multiple VPCs by using the most appropriate services based on requirements (for example, using VPC peering, Transit Gateway, PrivateLink)
- Using VPC sharing in a multi-account setup
- Managing IP overlaps by using different available services and options (for example, NAT, PrivateLink, Transit Gateway routing)

Domain 2: Network Implementation (26% of the exam content)

This domain accounts for 26% of the exam content.

Rubriques

- [Task 2.1: Implement routing and connectivity between on-premises networks and the Cloud](#)

- [Task 2.2: Implement routing and connectivity across multiple accounts, Regions, and VPCs to support different connectivity patterns](#)
- [Task 2.3: Implement complex hybrid and multi-account DNS architectures](#)
- [Task 2.4: Automate and configure network infrastructure](#)

Task 2.1: Implement routing and connectivity between on-premises networks and the Cloud

Knowledge of:

- Routing protocols (for example, static, dynamic)
- VPNs (for example, security, accelerated VPN)
- Layer 1 and types of hardware to use (for example, Letter of Authorization [LOA] documents, colocation facilities, Direct Connect)
- Layer 2 and layer 3 (for example, VLANs, IP addressing, gateways, routing, switching)
- Traffic management and SD-WAN (for example, Transit Gateway Connect)
- DNS (for example, conditional forwarding, hosted zones, resolvers)
- Security appliances (for example, firewalls)
- Load balancing (for example, layer 4 compared with layer 7, reverse proxies, layer 3)
- Infrastructure automation
- Organizations and Resource Access Manager (RAM) (for example, multi-account Transit Gateway, Direct Connect, Amazon VPC, Route 53)
- Test connectivity (for example, Route Analyzer, Reachability Analyzer)
- Networking services of VPCs

Skills in:

- Configuring the physical network requirements for hybrid connectivity solutions
- Configuring static or dynamic routing protocols to work with hybrid connectivity solutions
- Configuring existing on-premises networks to connect with the Cloud
- Configuring existing on-premises name resolution with the Cloud
- Configuring and implementing load balancing solutions

- Configuring network monitoring and logging for services
- Testing and validating connectivity between environments

Task 2.2: Implement routing and connectivity across multiple accounts, Regions, and VPCs to support different connectivity patterns

Knowledge of:

- Inter-VPC and multi-account connectivity (for example, VPC peering, Transit Gateway, VPN, third-party vendors, SD-WAN, multi-protocol label switching [MPLS])
- Private application connectivity (for example, PrivateLink)
- Methods of expanding networking connectivity (for example, Organizations, RAM)
- Host and service name resolution for applications and clients (for example, DNS)
- Infrastructure automation
- Authentication and authorization (for example, SAML, Active Directory)
- Security (for example, security groups, network ACLs, Network Firewall)
- Test connectivity (for example, Route Analyzer, Reachability Analyzer, tooling)

Skills in:

- Configuring network connectivity architectures by using services in a single-VPC or multi-VPC design (for example, DHCP, routing, security groups)
- Configuring hybrid connectivity with existing third-party vendor solutions
- Configuring a hub-and-spoke network architecture (for example, Transit Gateway, transit VPC)
- Configuring a DNS solution to make hybrid connectivity possible
- Implementing security between network boundaries
- Configuring network monitoring and logging by using solutions

Task 2.3: Implement complex hybrid and multi-account DNS architectures

Knowledge of:

- When to use private hosted zones and public hosted zones

- Methods to alter traffic management (for example, based on latency, geography, weighting)
- DNS delegation and forwarding (for example, conditional forwarding)
- Different DNS record types (for example, A, AAAA, TXT, pointer records, alias records)
- DNSSEC
- How to share DNS services between accounts (for example, RAM)
- Requirements and implementation options for outbound and inbound endpoints

Skills in:

- Configuring DNS zones and conditional forwarding
- Configuring traffic management by using DNS solutions
- Configuring DNS for hybrid networks
- Configuring appropriate DNS records
- Configuring DNSSEC on Route 53
- Configuring DNS within a centralized or distributed network architecture
- Configuring DNS monitoring and logging on Route 53

Task 2.4: Automate and configure network infrastructure

Knowledge of:

- Infrastructure as code (IaC) (for example, Cloud Development Kit [CDK], CloudFormation, CLI, SDK, APIs)
- Event-driven network automation
- Common problems of using hardcoded instructions in IaC templates when provisioning cloud networking resources

Skills in:

- Creating and managing repeatable network configurations
- Integrating event-driven networking functions
- Integrating hybrid network automation options with native IaC

- Eliminating risk and achieving efficiency in a cloud networking environment while maintaining the lowest possible cost
- Automating the process of optimizing cloud network resources with IaC

Domain 3: Network Management and Operation (20% of the exam content)

This domain accounts for 20% of the exam content.

Rubriques

- [Task 3.1: Maintain routing and connectivity on and hybrid networks](#)
- [Task 3.2: Monitor and analyze network traffic to troubleshoot and optimize connectivity patterns](#)
- [Task 3.3: Optimize networks for performance, reliability, and cost-effectiveness](#)

Task 3.1: Maintain routing and connectivity on and hybrid networks

Knowledge of:

- Industry-standard routing protocols that are used in hybrid networks (for example, BGP over Direct Connect)
- Connectivity methods for and hybrid networks (for example, Direct Connect gateway, Transit Gateway, VIFs)
- How limits and quotas affect networking services (for example, bandwidth limits, route limits)
- Available private and public access methods for custom services (for example, PrivateLink, VPC peering)
- Available inter-Regional and intra-Regional communication patterns

Skills in:

- Managing routing protocols for and hybrid connectivity options (for example, over a Direct Connect connection, VPN)
- Maintaining private access to custom services (for example, PrivateLink, VPC peering)
- Using route tables to direct traffic appropriately (for example, automatic propagation, BGP)
- Setting up private access or public access to services (for example, Direct Connect, VPN)
- Optimizing routing over dynamic and static routing protocols (for example, summarizing routes, CIDR overlap)

Task 3.2: Monitor and analyze network traffic to troubleshoot and optimize connectivity patterns

Knowledge of:

- Network performance metrics and reachability constraints (for example, routing, packet size)
- Appropriate logs and metrics to assess network performance and reachability issues (for example, packet loss)
- Tools to collect and analyze logs and metrics (for example, CloudWatch, VPC Flow Logs, VPC Traffic Mirroring)
- Tools to analyze routing patterns and issues (for example, Reachability Analyzer, Transit Gateway Network Manager)

Skills in:

- Analyzing tool output to assess network performance and troubleshoot connectivity (for example, VPC Flow Logs, Amazon CloudWatch Logs)
- Mapping or understanding network topology (for example, Transit Gateway Network Manager)
- Analyzing packets to identify issues in packet shaping (for example, VPC Traffic Mirroring)
- Troubleshooting connectivity issues that are caused by network misconfiguration (for example, Reachability Analyzer)
- Verifying that a network configuration meets network design requirements (for example, Reachability Analyzer)
- Automating the verification of connectivity intent as a network configuration changes (for example, Reachability Analyzer)
- Troubleshooting packet size mismatches in a VPC to restore network connectivity

Task 3.3: Optimize networks for performance, reliability, and cost-effectiveness

Knowledge of:

- Situations in which a VPC peer or a transit gateway are appropriate
- Different methods to reduce bandwidth utilization (for example, unicast compared with multicast, CloudFront)
- Cost-effective connectivity options for data transfer between a VPC and on-premises environments

- Different types of network interfaces on
- High-availability features in Route 53 (for example, DNS load balancing using health checks with latency and weighted record sets)
- Availability of options from Route 53 that provide reliability
- Load balancing and traffic distribution patterns
- VPC subnet optimization
- Frame size optimization for bandwidth across different connection types

Skills in:

- Optimizing for network throughput
- Selecting the right network interface for the best performance (for example, elastic network interface, Elastic Network Adapter [ENA], Elastic Fabric Adapter [EFA])
- Choosing between VPC peering, proxy patterns, or a transit gateway connection based on analysis of the network requirements provided
- Implementing a solution on an appropriate network connectivity service (for example, VPC peering, Transit Gateway, VPN connection) to meet network requirements
- Implementing a multicast capability within a VPC and on-premises environments
- Creating Route 53 public hosted zones and private hosted zones and records to optimize application availability (for example, private zonal DNS entry to route traffic to multiple Availability Zones)
- Updating and optimizing subnets for auto scaling configurations to support increased application load
- Updating and optimizing subnets to prevent the depletion of available IP addresses within a VPC (for example, secondary CIDR)
- Configuring jumbo frame support across connection types
- Optimizing network connectivity by using Global Accelerator to improve network performance and application availability

Domain 4: Network Security, Compliance, and Governance (24% of the exam content)

This domain accounts for 24% of the exam content.

Rubriques

- [Task 4.1: Implement and maintain network features to meet security and compliance needs and requirements](#)
- [Task 4.2: Validate and audit security by using network monitoring and logging services](#)
- [Task 4.3: Implement and maintain confidentiality of data and communications of the network](#)

Task 4.1: Implement and maintain network features to meet security and compliance needs and requirements

Knowledge of:

- Different threat models based on application architecture
- Common security threats
- Mechanisms to secure different application flows
- network architecture that meets security and compliance requirements

Skills in:

- Securing inbound traffic flows into (for example, WAF, Shield, Network Firewall)
- Securing outbound traffic flows from (for example, Network Firewall, proxies, Gateway Load Balancers)
- Securing inter-VPC traffic within an account or across multiple accounts (for example, security groups, network ACLs, VPC endpoint policies)
- Implementing an network architecture to meet security and compliance requirements (for example, untrusted network, perimeter VPC, three-tier architecture)
- Developing a threat model and identifying appropriate mitigation strategies for a given network architecture
- Testing compliance with the initial requirements (for example, failover test, resiliency)
- Automating security incident reporting and alerting using

Task 4.2: Validate and audit security by using network monitoring and logging services

Knowledge of:

- Network monitoring and logging services that are available in (for example, CloudWatch, CloudTrail, VPC Traffic Mirroring, VPC Flow Logs, Transit Gateway Network Manager)
- Alert mechanisms (for example, CloudWatch alarms)
- Log creation in different services (for example, VPC flow logs, load balancer access logs, CloudFront access logs)
- Log delivery mechanisms (for example, Amazon Kinesis, Route 53, CloudWatch)
- Mechanisms to audit network security configurations (for example, security groups, Firewall Manager, Trusted Advisor)

Skills in:

- Creating and analyzing a VPC flow log (including base and extended fields of flow logs)
- Creating and analyzing network traffic mirroring (for example, using VPC Traffic Mirroring)
- Implementing automated alarms by using CloudWatch
- Implementing customized metrics by using CloudWatch
- Correlating and analyzing information across single or multiple log sources
- Implementing log delivery solutions
- Implementing a network audit strategy across single or multiple network services and accounts (for example, Firewall Manager, security groups, network ACLs)

Task 4.3: Implement and maintain confidentiality of data and communications of the network

Knowledge of:

- Network encryption options that are available on
- VPN connectivity over Direct Connect
- Encryption methods for data in transit (for example, IPsec)
- Network encryption under the shared responsibility model
- Security methods for DNS communications (for example, DNSSEC)

Skills in:

- Implementing network encryption methods to meet application compliance requirements (for example, IPsec, TLS)
- Implementing encryption solutions to secure data in transit (for example, CloudFront, Application Load Balancers and Network Load Balancers, VPN over Direct Connect, managed databases, Amazon S3, custom solutions on Amazon EC2, Transit Gateway)
- Implementing a certificate management solution by using a certificate authority (for example, ACM, Private Certificate Authority [ACM PCA])
- Implementing secure DNS communications

In-Scope AWS Services

The following list contains AWS services and features that are in scope for the AWS Certified Advanced Networking - Specialty (ANS-C01) exam. This list is non-exhaustive and is subject to change. AWS offerings appear in categories that align with the offerings' primary functions.

Rubriques

- [Application Integration](#)
- [Compute](#)
- [Containers](#)
- [Cost Management](#)
- [Front-End Web and Mobile](#)
- [Management and Governance](#)
- [Networking and Content Delivery](#)
- [Security, Identity, and Compliance](#)
- [Storage](#)

Application Integration

- Amazon EventBridge
- Amazon Simple Notification Service (Amazon SNS)
- Amazon Simple Queue Service (Amazon SQS)

Compute

- Amazon EC2
- Amazon EC2 Auto Scaling
- AWS Lambda

Containers

- Amazon Elastic Container Registry (Amazon ECR)
- Amazon Elastic Container Service (Amazon ECS)
- Amazon Elastic Kubernetes Service (Amazon EKS)
- AWS Fargate

Cost Management

- AWS Cost Explorer

Front-End Web and Mobile

- Amazon API Gateway

Management and Governance

- AWS Auto Scaling
- AWS CLI
- AWS CloudFormation
- AWS Cloud Development Kit (AWS CDK)
- AWS CloudTrail
- Amazon CloudWatch
- AWS Config
- AWS Control Tower
- AWS Organizations
- AWS Resource Access Manager (AWS RAM)

- AWS Systems Manager
- AWS Trusted Advisor

Networking and Content Delivery

- Amazon CloudFront
- AWS Direct Connect
- Elastic Load Balancing (ELB)
- AWS Global Accelerator
- AWS Network Firewall
- AWS PrivateLink
- Amazon Route 53
- AWS Site-to-Site VPN
- AWS Transit Gateway
- Amazon VPC
- AWS VPN

Security, Identity, and Compliance

- AWS Certificate Manager (ACM)
- AWS Firewall Manager
- AWS Identity and Access Management (IAM)
- AWS Key Management Service (AWS KMS)
- AWS Private Certificate Authority (AWS Private CA)
- AWS Security Hub
- AWS Shield
- AWS WAF

Storage

- Amazon Elastic Block Store (Amazon EBS)

- Amazon Elastic File System (Amazon EFS)
- Amazon FSx
- Amazon S3

Out-of-Scope AWS Services

The following list contains AWS services and features that are out of scope for the AWS Certified Advanced Networking - Specialty (ANS-C01) exam. This list is non-exhaustive and is subject to change. AWS offerings that are entirely unrelated to the target job roles for the exam are excluded from this list.

Rubriques

- [Analytics](#)
- [Business Applications](#)
- [Database](#)
- [Developer Tools](#)
- [End User Computing](#)
- [Machine Learning](#)
- [Media Services](#)
- [Migration and Transfer](#)

Analytics

- Amazon Athena
- Amazon EMR
- Amazon Kinesis
- Amazon OpenSearch Service
- Amazon QuickSight

Business Applications

- Amazon Connect
- Amazon Simple Email Service (Amazon SES)

- Amazon WorkMail

Database

- Amazon Aurora
- Amazon DynamoDB
- Amazon ElastiCache
- Amazon RDS
- Amazon Redshift

Developer Tools

- AWS CodeBuild
- AWS CodeCommit
- AWS CodeDeploy
- AWS CodePipeline
- AWS X-Ray

End User Computing

- Amazon AppStream 2.0
- Amazon WorkSpaces

Machine Learning

- Amazon Comprehend
- Amazon Lex
- Amazon Polly
- Amazon Rekognition
- Amazon SageMaker
- Amazon Textract
- Amazon Transcribe
- Amazon Translate

Media Services

- Amazon Elastic Transcoder
- Amazon Kinesis Video Streams

Migration and Transfer

- AWS Database Migration Service (AWS DMS)
- AWS DataSync
- AWS Migration Hub
- AWS Snow Family
- AWS Transfer Family

Document history for the AWS Certification User Guide

The following table describes the documentation releases for AWS Certification.

Modification	Description	Date
Initial release	Initial release of the AWS Certification User Guide	October 18, 2025